Visual Information and AI Divide: Evidence from Corporate Executive Presentations*

Sean Cao Yichen Cheng Meng Wang Yusen Xia

Baozhong Yang

This Draft: June 2024

ABSTRACT

This paper constructs and studies a novel data set comprised of corporate executive presentations, which provide unique visual information about firms' products and operations. We explore the value of visual information in presentations and examine how market participants respond to such information. We extract visual features from presentation images with large image models and find not all images have the same implications for firm value. Forward-looking operational information is associated with higher short-term abnormal returns and long-term operational performance while other types of images are not. We also examine whether the rise of alternative big data and AI creates a potential *AI divide* among market participants. We find AI-equipped institutional investors respond strongly to visual signals, whereas retail investors and traditional institutional investors face marginalization on the playing field in the age of AI and big data.

JEL Classification: D83, G12, G14, G23, G30

*The authors have benefited from discussions with Vikas Agarwal, Peter Cziraki, Kuntara Pukthuanthong (discussant), Gang Li (discussant), Russell Jame, Danqing Mei, Dan Su, and Xingtan Zhang, and comments and suggestions from participants in seminars and conferences at 2024 Financial Accounting and Reporting Section (FARS) Midyear Meeting, 2024 Hawaii Accounting Research Conference, The 2023 Research Conference on Capital Market Research in the Era of AI, CICF 2023, 2023 Hong Kong Conference for Fintech, AI, and Big Data in Business, AAAI 23, 16th NYCU Finance Conference, CKGSB, Georgia State University, University of Maryland, HKUST, Xiamen University, Nankai University, and Renmin University. We gratefully acknowledge the support from the Smith AI Initiative for Capital Market Research at the University of Maryland. Cao is with the Robert H. Smith School of Business University of Maryland College Park, Email: scao824@umd.edu. Cheng, Wang, Xia, and Yang are with the J. Mack Robinson College of Business, Georgia State University; Cheng, Email: ycheng11@gsu.edu; Wang, Email: mwang36@gsu.edu; Xia, Email: ysxia@gsu.edu; (Contact author) Yang, Email: bzyang@gsu.edu.

Key Words: Machine Learning, Artificial Intelligence, FinTech, Corporate Presentation, Image Analysis, Textual Analysis, Information and Market Efficiency, Institutional Investors, Investment

1 Introduction

Information disclosed by corporations or media plays a fundamental role in the financial markets. Investors rely on such information to form expectations and shape asset prices, and may obtain it from various sources, including financial statement releases, conference calls, annual reports, and the news media. In this paper, we explore information from an relatively unexplored firm information event: corporate executive presentations, which has become a prevalent information channel for corporate managers and capital market participants. Corporate executive presentations are a unique setting to study visual information in two ways. First, since CEOs must deliver live presentations within a given time limit, their slides tend to include a large amount of visual and graphic elements with minimal verbal explanation. This makes the visual information largely independent of narrative explanations. Second, in contrast with other corporate disclosures, executive presentations provide an abundance of visual information about a firm's product designs and operation plans.

Due in large part to its unique capacity for analyzing alternative data such as images on a massive scale, artificial intelligence is now being deployed with increasing frequency both by the financial services industry[1] and in academia.[2] Motivated by the potential that emergent AI technologies hold for the effective processing of visual data from corporate presentations, this paper addresses two central research questions. First, we study whether and how state-of-the-art machine learning techniques and large image models can extract valuable information from visual data and help investors understand corporate business operations and future performance. Second, utilizing this setting, we examine whether AI technologies create disparities between market participants by marginalizing retail investors and traditional institutional investors lacking AI resources.

We first construct a comprehensive data set of corporate presentation slides from multi-

[1] According to a survey report by BarclayHedge, over half of hedge fund respondents (56%) used AI to inform investment decisions in 2018. The report is available at `https://www.barclayhedge.com/insider/majority-of-hedge-fund-pros-use-ai-machine-learning-in-investment-strategies`. See also a survey of the recent AI literature in financial economics at the end of this introduction.

[2] See for example the survey Cong, Liang, Yang, and Zhang (2020).

ple data sources including Bloomberg News and corporate websites from 2005 to 2018. Our sample includes multiple types of executive presentation events, including non-deal road shows (e.g., Bradley, Jame, and Williams, 2021; Ellis, Gerken, and Jame, 2020), IPO road shows (e.g., Blankespoor et al., 2017, 2020b), broker-hosted investor conferences (e.g., Green, Jame, Markov, and Susai, 2014), capital market day events, and others. The final sample used for our analysis includes 17,277 corporate presentation slide decks associated with an average of 1,023 unique firms per year and consisting of 464,765 slide pages. Although quantitative financial information is disclosed to the market in various ways (e.g., earnings announcements, SEC filings), corporate presentations provide a unique opportunity for companies to visually showcase the details of their business operations to potential investors.[3] Furthermore, managers can utilize graphics to better present forward-looking information that can affect future cash flows and the value of the company. To reflect the nature of information conveyed by corporate images, we use a deep learning algorithm to classify them into three categories: 1) *Operations Forward*: images in this category provide forward-looking operational information, including future products, blueprints, and development plans; 2) *Operations Summary*: images in this category present information about existing products or services; 3) *Others*, which includes graphs of financial and quantitative information and generic images.

In our classification process, we adopt a deep learning model specifically tailored for image recognition: convolutional neural networks (CNN). Furthermore, we utilize the transfer learning method (e.g., Pratt, 1993) to improve prediction accuracy and reduce the size of the training sample. Transfer learning utilizes and fine-tunes existing deep neural network models that have been pre-trained with large labelled image data sets.[4]

We first examine the value of visual information extracted from corporate presentation

[3]While companies also disclose graphics in 10-K reports, the images therein are mostly plots of quantitative data. Corporate presentations allow us to focus on images that provide information about corporate operations and products.

[4]In particular, we leverage the pre-trained VGG16 (Simonyan and Zisserman, 2014) model using the Google ImageNet data. We provide details about CNN and transfer learning models in Appendix A.

images by testing whether visual information can predict stock returns. We hypothesize that *Operations Forward* contains new, valuable information for investment, whereas the information in *Operations Summary* will, for the most part, have already been incorporated into stock prices. As expected, we find that *Operations Forward* is associated with significant and positive short-term cumulative announcement returns ([-3, 3] relative to the presentation date), while other types of visual information are not. A one-standard-deviation increase in *Operations Forward* is associated with an increase in abnormal returns of approximately 14 basis points around the presentation date. To partial out potential endogenous impacts of textual information embedded in slides, we re-run similar training process using textual contents embedded in slides alone and construct measures for textual operational information. We find that the predictive power of *Operations Forward* remains robust after controlling for textual operational information.

Information can affect stock prices via two channels: the discount rate and future cash flows (Campbell and Shiller, 1988). Given that corporate disclosure of prospective operations is unlikely to be directly related to discount rates, we expect it to influence stock prices through the cash flow channel. Consistent with this hypothesis, we find that *Operations Forward* is positively associated with firms' sales and earnings in the fourth quarter and the second year following the presentation. Interestingly, *Operations Summary* is only significantly correlated with sales and earnings in the next quarter, but not in the long run, indicating that it contains mostly stale and short-term information.

We next explore different market participants' responses to the visual information contained in corporate presentations. Drawing on recent studies showing that unequal access to alternative data and capacity to analyze data at scale increases information asymmetry among market participants (e.g., Katona et al., 2018; Zhu, 2019), we hypothesize that the ability to process and extract unique information from unstructured data provides information advantages to institutions that have adopted AI technologies. Specifically, since extracting visual information from slides on a large scale requires deep learning capabilities, we expect that AI-equipped financial institutions are more likely to trade on visual information, compared with other institutions and retail investors. We measure the extent of AI adoption by institutional investors by calculating the cumulative job postings in the BurningGlass database.[5] Consistent with our hypothesis, we find that institutional investors with high AI investment trade more around the presentation date when corporate presentations contain more forward-looking visual information related to operations (higher *Operations Forward*). In contrast, trades by other institutional investors and retail investors are not sensitive to the visual information contained in CEO presentations.[6]

If AI-equipped institutions are more able to process visual information, then stock prices should incorporate such information more quickly when these institutions are present. To test this hypothesis, we separate stocks into two groups, based on the level of ownership by AI-equipped institutions. We find that *Operations Forward* is only associated with significant abnormal returns around corporate presentations when the stock has a high proportion of AI-equipped institutional ownership, suggesting that AI-guided institutional trades are what impound the visual operations information into prices.

This study contributes to several strands of literature. First, our work extends the growing literature on the impact of AI in the financial services industry. For example, Abis (2020) examines how quantitative investment strategies influence mutual fund performance. Grennan and Michaely (2020) study how analysts perform and adjust in response to the advent of AI-processed recommendations in the markets. Abis and Veldkamp (2020) explore the change in labor shares in the financial industry driven by new data management and AI-related jobs. Coleman, Merkley, and Pacelli (2021) compare the performance of robot analysts from FinTech companies to that of human analysts. Cao, Jiang, Yang, and Zhang (2020) study how firms modify their disclosure practices in response to the presence of AI readership in financial markets. Our paper provides direct evidence of the value of visual

[5]Our calculation follows a similar approach to those in Acemoglu et al. (2020), Abis and Veldkamp (2020), and Babina et al. (2021).

[6]Retail trades are calculated following Boehmer et al. (2021).

information extracted by AI. We show that AI technologies can create information and trading advantages for adopting institutions, potentially generating an *AI divide* among investors.

Second, we contribute to the emerging literature on the value of visual content in corporate disclosures and news media. Several recent studies document the importance of visual content in predicting many outcomes. Deng, Gao, Hu, and Zhou (2020) find that firms that increase their use of graphics experience abnormal returns over the next six months. Ben-Rephael, Ronen, Ronen, and Zhou (2021) measure the visual readability of annual reports and find that images serve as an important source of information by reinforcing firms' textual narratives. Christensen, Fronk, Lee, and Nelson (2021) find a substantial increase in the disclosure of infographics in the last decade. Obaid and Pukthuanthong (2022) construct a sentiment measure using photos in news articles, demonstrating that negative sentiment predicts market return reversals and trading volume. In this paper, we construct a new data set, comprised of corporate presentations, which, since such information is more abundant and salient therein, provides a unique setting to study visual information about business operations and products. Different from the above studies that study images and their relation to texts, we study an information environment where the textual narrative is more limited and focus on comparing different types of visual information contained in images. Interestingly, we find that only certain types of visual information (forward-looking, operational information) impact prices, and that investors have differential abilities to utilize such information.

Finally, our paper contributes to a growing body of studies that apply machine learning and deep learning techniques to financial issues by, for example, using it to predict asset prices (Gu, Kelly, and Xiu, 2020; Brogaard and Zareei, 2022), manage portfolios (Chen, Pelger, and Zhu, 2022; Cong, Tang, Wang, and Zhang, 2022), forecast earnings (van Binsbergen, Han, and Lopez-Lira, 2022; Cao and You, 2021), and analyze audio and video information (Mayew and Venkatachalam, 2012; Hu and Ma, 2021). To complement this literature, our paper explores transfer learning techniques and CNN models, ultimately developing a framework for future studies of visual information in important corporate disclosures.

The rest of the paper is structured as follows. Section 2 describes our data, model, and the method used to construct key variables. Section 3 and Section 4 present empirical results. Section 5 concludes.

2 Visual Information and Machine Learning

2.1 Corporate Presentations Data

We collect a comprehensive dataset of corporate presentation slides from multiple data sources, including Bloomberg News and corporate websites, generated in the period from 2005 to 2018. The corporate presentations in our sample include all types of public executive presentations that are hosted either by the corporations themselves (e.g., corporate conferences) or by third parties (e.g., investment banks). The vast majority of presentations in our sample are held externally and are non-road shows. Our initial sample contains 51,879 corporate presentations (slide decks) and consists of 1,467,979 slide pages that include both graphical and textual content. After merging with CRSP/Compustat[7], our final sample consists of 17,277 corporate presentation slide decks associated with an average of 1,023 unique public US firms per year and includes 464,765 slide pages. Each slide page is converted to an image as the input for the machine learning algorithms. We also use a Python algorithm to extract the textual information contained in the slides.

Visual and graphic information has been difficult to analyze due to its unstructured and high-dimensional nature. An image may contain tens of thousands of pixels, each with millions of possible colors, that form complex patterns and objects. Recent advances in machine learning and AI, however, have made it possible to construct image recognition algorithms with capabilities comparable to humans. In this section, we describe how we

[7] All common variables used are defined as in Appendix A.

apply deep learning to extract key features of firms' operations from corporate presentations.

2.2 Image Labeling and Classification

As the first step, we manually review and classify (label) a random sub-sample of images into several different categories, providing a training sample for the machine learning algorithms we later employ. We classify each image into one of three categories: *Operations Forward*, *Operations Summary*, and *Others*. Our labeling criteria are as follows:

1) *Operations Forward*: If a slide image mainly provides information regarding a firm's future operations, constructions, products, or programs, we classify it as an *Operations Forward* image. This would include, for example, images that present new product designs, blueprints, or future business plans. We present examples of images from each category in Appendix C.

2) *Operations Summary*: If an image mostly presents information about a firm's existing operations and products, we label it as an *Operations Summary* image. For example, such images might show existing products, established factories, or current business strategies.

3) *Others*: If an image is a plot or chart (usually based on financial information that is also disclosed elsewhere), if the image contains mostly text, or if it contains information not directly relevant to the firm's operations (e.g., generic images, logos), we label it as *Others*.

To minimize human error in the labeling process, we cross-validate and require a consensus on the classification by at least three graduate research assistants. We use a two-step bootstrapping process to construct the training sample. We first label an initial random sample of 3,000 images. We then use this initial sample to train our machine learning model (described in the next section) and make initial predictions on how the remaining images will be classified. We then select a final training sample of 20,000 pre-classified images with a balanced number of images in each category, and manually classify images in the sample.

2.3 Deep Learning Model

Over time, a diverse array of applications have been found for various machine learning models, including random forests, gradient boosting, and neural networks, have found use in a wide range of applications (e.g., James, Witten, Hastie, and Tibshirani, 2013). Image recognition, however, has long posed a significant problem for deep learning. Hence Google's development of ImageNet (Li et al., 2009), which is able to perform on par with humans, represented a major milestone. The primary deep learning model employed by ImageNet and other leading image recognition algorithms is the Convolutional Neural Network (CNN). The CNN is a multi-layer neural network in which the lower layers capture finer details while the higher layers extract high-level information, such as objects in the image. We include a detailed discussion of the structure and intuition of CNN in Appendix B.

Still, recognizing business-related images is challenging because there are no ready-made models for this purpose, and training a CNN model usually requires a large training dataset. Therefore, we utilize an advanced machine learning technique called *transfer learning* (Pratt, 1993; Rajat et al., 2006) to build our own deep learning model based on pre-trained CNN models and then train the model with our business image sample. Specifically, we first build a neural network on top of a pre-trained CNN neural network from a state-of-the-art image recognition model VGG16 (Simonyan and Zisserman, 2014). We then keep the parameters of CNN layers fixed and fine-tune the model with our training sample. The resulting model is what we call the *Transfer CNN* model. Transfer learning allows us to take advantage of existing CNN models trained with very large datasets while adapting the model to our specific business problem. We also employ the model that utilizes both image and text information from presentations (*Transfer CNN + Text*) and the transfer learning technique. More details of our models are described in Appendix B.

2.4 Model Performance

We consider and train four different model architectures: 1) A CNN model that starts from scratch (*CNN*); 2) A deep learning model that processes both images and text (*CNN + Text*); 3) A transfer learning model that relies on a pre-trained CNN model to process images (*Transfer CNN*); and 4) A transfer learning model that processes both images and text (*Transfer CNN + Text*).[8]

[Insert Table 1 Here]

We use four measures to evaluate the models' out-of-sample performance. *Accuracy* is the ratio of correct predictions to total observations. *Precision* is the ratio of true positives to the sum of true positives and false positives. *Recall* is the ratio of true positives to the sum of true positives and false negatives. We calculate *Precision* and *Recall* for each category and then take an average across all three categories. *F1 score* is the harmonic mean of *Precision* and *Recall*. Among the four architectures, *Transfer CNN* and *Transfer CNN + Text* have the best performance in terms of F1 score and accuracy. *Transfer CNN + Text* outperforms other models with an accuracy of 80.0% and F1 score of 79.4%. We therefore use it as our main model in the paper.

At the same time, it should be noted that the difference in performance between the *Transfer CNN* and *Transfer CNN + Text* is small, amounting to only 1.1% for the F1 score and 3% for accuracy. This suggests that the vast majority of the prediction power of our model comes from visual information rather than textual information. Intuitively, the textual information in the executive presentation slides tends to be brief and is often used only to supplement the visual information.[9]

[8]Textual information might also be useful in helping classifying forward-looking v.s. backward-looking information. Keywords potentially indicate forward-looking information include “projected”, “prospect”, “expected”, “guidance”, “target”, “potential”, “forecast”, “estimated”, “will”, “plan”, “outlook”, “opportunity”, “explore”, “expand”, “develop”, “future, “pipeline”, “new”, “proposal”, “schedule”, “next”, “update”, “forward”, “prepare”.

[9]We were able to obtain presentation transcripts for a small subset of our sample from Bloomberg. Due to the limited size, we cannot conduct a comprehensive analysis using the verbal information from transcripts.

3 Variable Construction

3.1 Operations Information from Images

We use the trained *Transfer CNN + Text* model to obtain a final classification for the entire sample of 464,765 slide images in 17,277 corporate presentations slide decks. For each presentation, we then calculate the fraction of slide pages in each category to define our main variables of operation information. Specifically, for each presentation i, we define

$$Operations\ Forward_i = \frac{\#\ of\ Operations\ Forward\ Slides_i}{Total\ Number\ of\ Slides_i}, \tag{1}$$

$$Operations\ Summary_i = \frac{\#\ of\ Operations\ Summary\ Slides_i}{Total\ Number\ of\ Slides_i}. \tag{2}$$

These two variables represent the extent of visual information about forward-looking and backward-looking operations in the executive presentations, respectively. They form the main independent variables in our subsequent analyses.

3.2 AI Institutional Ownership

In this section, we briefly describe our methodology for identifying AI-equipped financial institutions and their stock ownership. We also describe our methodology in greater detail in Appendix D. First, we estimate a financial institution's AI-related labor stock based on job posting data from Burning Glass. This methodology is similar to that of Acemoglu et al. (2020), Abis and Veldkamp (2020), and Babina et al. (2021). Intuitively, a given financial institution's AI labor stock in a certain year is calculated based on the time series of the number of AI-related job postings, and the estimated hiring/separation rate in the financial services sector. We then classify financial institutions into AI-equipped and non-AI-equipped institutions based on their estimated AI-related labor stock. Finally, we calculate stock-level ownership by AI-equipped institutions.

A given financial institution i in month t is defined as an *AI Institution* if its AI labor stock in the preceding year is among the top 30% of the sample. We aggregate 13F stock holdings of AI institutions to construct stock-level AI institutional ownership or *AI Inst. Ownership*. Specifically, for a stock-month (i, t),

$$AI\ Inst.\ Ownership_{i,t} = \frac{AI\ Institutional\ Shares_{i,t}}{Shares\ Outstanding_{i,t}}, \tag{3}$$

where $AI\ Institutional\ Shares_{i,t}$ is the aggregate 13F institutional holdings of stock i at time t (if t is not at quarter end, the previous quarter-end holdings are used), and $Shares\ Outstanding_{i,t}$ is the total shares outstanding of stock i at the end of month t.

3.3 Retail Trades

We follow Boehmer et al. (2021) and use Trade and Quote (TAQ) data to calculate marketable retail order imbalance. After identifying marketable retail buy and sell transactions following Boehmer et al. (2021), we aggregate transactions at the stock-day level and calculate the marketable retail trading measures as follows, for stock-day (i, t),

$$Retail\ Order\ Imbalance\ (Shares)_{i,t} = \frac{mrbvol_{i,t} - mrsvol_{i,t}}{mrbvol_{i,t} + mrsvol_{i,t}}, \tag{4}$$

$$Retail\ Order\ Imbalance\ (Trades)_{i,t} = \frac{mrbtrd_{i,t} - mrstrd_{i,t}}{mrbtrd_{i,t} + mrstrd_{i,t}}, \tag{5}$$

where $mrbvol_{i,t}, mrsvol_{i,t}$ are the daily buy and sell share volume, respectively, and $mrbtrd_{i,t}$, $mrstrd_{i,t}$ are the daily number of retail buy and sell trades.

3.4 Textual Variables

We employ the Loughran and McDonald (2011) dictionaries and follow their methodology to construct textual sentiment and tone measures for presentations: *Positive Sentiment*, *Negative Sentiment*, *Uncertainty*, and *Constrained*. For example, *Positive Sentiment* is the ratio of words in the positive LM dictionary to the total number of words in the presentation

(*Text Length*). Other measures are similarly defined. We also construct a comprehensive dictionary of forward-looking words by combining the dictionaries from Li (2010), Muslu et al. (2015), Bozanic et al. (2018), and Grewal et al. (2019) and compute *Forward Looking* as the ratio of the number of forward-looking words to the text length of the slides. We also build a neural network based only on textual information from the corporate presentation slides to predict the slide categories and construct the *Textual Operations Forward* and *Textual Operations Summary* variables, employing methods similar to those used to construct the visual information variables. Given that the focus of this study is on visual information, we include textual variables as controls in our analyses.

3.5 Other Firm Characteristics

We also consider a number of other firm characteristics as control variables in our tests. *Size* is defined as the natural logarithm of the market capitalization. *Book-to-Market* is the total assets over market cap plus the book value of liabilities. *Turnover* is the monthly average of the ratio of trading volume to shares out- standing, multiplied by twelve. *Inst. Ownership* is the ratio of the total shares of institutional ownership to shares outstanding.

3.6 Summary Statistics

Table 2 reports the summary statistics of the variables. On average, 3.6% of the slides in a presentation are *Operations Forward* slides and 11% are *Operations Summary* slides. There is substantial cross-sectional variation in these variables, as the standard deviation of *Operations Forward* is 5.8%, and *Operations Summary* is 12.2%.

[Insert Table 2 Here]

We further identify the industries with the highest concentration of *Operations Forward* information in Table 3. Panel A of Table 3 displays the top 10 two-digit SIC industries arranged by the fraction of presentations that contain *Operations Forward* slides. Panel B shows the top 10 industries by the average *Operations Forward* in the industry. The top industries include capital-intensive industries such as the Metal, Mining, Oil & Gas, Transportation, and Construction industries, for which future investment plans are important, as well as various Consumer Product industries (e.g., Food, Apparel, and Leather), for which new product designs could be influential.

[Insert Table 3 Here]

Figure 2 shows the time series of the number and percentage of presentations with *Operations Forward* information in our sample from 2006 to 2018. The number of presentations with Operations Forward slides increased steadily over our sample period, due in part to the general increase in the total number of presentations. Interestingly, the ratio of presentations with *Operations Forward* images also increases over time, from 35% in 2006 to 47% in 2018. The increasing propensity of firms for including *Operations Forward* visual information in their public presentations may be due to the increasing demand for new, rich information by investors in the era of AI and big data.

[Insert Figure 2 Here]

Figure 3 compares the distribution of presentation dates by month of the year with those of 10-K and 10-Q filings. Unlike 10-Ks and 10-Qs, which are clustered in common fiscal reporting months (primarily February, May, August, and November), corporate presentation events are relatively evenly distributed, indicating that they provide a continuous information flow throughout the year.

[Insert Figure 3 Here]

4 Visual Information and Capital Markets

4.1 Does Visual Information Matter to the Market?

In this section, we explore whether visual information from corporate presentations offers new insights for investors. The literature has shown that textual disclosures, such as 10-K filings and conference call transcripts, provide new information to the market (e.g. Tetlock, 2007; Li, 2010; Loughran and McDonald, 2011; Jiang, Lee, Martin, and Zhou, 2019). Given that visual information about corporate operations extracted from executive presentations is unique and not contained in common corporate disclosures, we hypothesize that they offer investors value-relevant information. Furthermore, since *Operations Summary* information is largely known to investors, while *Operations Forward* information is forward-looking and can be particularly helpful in forming expectations of future firm value, we expect *Operations Forward* to provide new information to investors. To the extent that managers enjoy discretion over which information is disclosed in corporate presentations, they are more likely to highlight positive outlooks on operations and omit negative prospects when possible. Therefore, we expect *Operations Forward* to be associated with positive stock reactions.

4.1.1 Visual Operational Information and Announcement Returns

We examine the market's reaction to visual information in presentations with the following regression for presentation-firm-date (i, j, t),

$$\begin{aligned} CAR(t_1, t_2)_{i,j,t} = \beta_0 + \beta_1 Operations\ Forward_{i,t} + \beta_2 Operations\ Summary_{i,t} \\ + \gamma_1 Textual\ Controls_{i,t} + \delta Firm\ Controls_{j,t} + \alpha_{SIC2} + \mu_t + \varepsilon_{i,j,t}, \end{aligned} \tag{6}$$

where $CAR(t_1, t_2)_{i,j,t}$ is the cumulative abnormal returns adjusted for the Fama-French three-factor model during the window $[t_1, t_2]$ around the presentation date t. We include a set of firm-level control variables such as *Size*, *Book-to-Market*, *Turnover*, *Nasdaq*, and *Institutional Ownership*. Textual control variables include *Positive Sentiment*, *Negative Sentiment*, *Uncer-*

tainty, *Constrained*, and *Forward Looking*. ML Textual Variables include the ML text-based variables *Textual Operations Forward* and *Textual Operations Summary*. We also include year and industry fixed effects.

[Insert Table 4 Here]

Table 4 reports the results. Consistent with our hypotheses, columns (1) and (2) show that cumulative abnormal returns in window $[-3, 3]$ are significantly and positively correlated with forward-looking visual operational information, but insignificantly correlated with backward operational information. In particular, column (2) shows that, in the model including control variables and industry and year fixed effects, a one-standard-deviation increase (0.058) in *Operations Forward* is associated with an increase of 14 basis points in $CAR(-3, 3)$ around the presentation date, or 5% annualized abnormal returns. Columns (3) and (4) show that the relationship between visual information and announcement returns holds for the shorter window $[-2, 2]$. Columns (5) and (6) reveal that $CAR(4, 13)$ is not significantly correlated with either *Operations Forward* or *Operations Summary*.

Overall, the results support the hypothesis that visual information about firm operations gleaned from corporate presentations provides new signals to investors. Such signals are concentrated in forward-looking visual information and are quickly incorporated into stock prices.

4.1.2 Textual Information in Slides

A natural question arising from our study is whether the visual information extracted from corporate presentation slides provides incremental value or merely duplicates the textual information contained within these slides. To test these possibilities, we re-run the same training process described in Section 2 but restrict inputs to textual contents from slides only.[10] We then use the trained model to classify slide pages into corresponding categories, and construct variables *Textual Operations Forward* and *Textual Operations Summary*, which

[10]Specifically, we train and fine-tune a fully connected neural network model for text classification.

are defined as the ratio of the number of slides with the corresponding classification to the total number of slides. Next, we analyze whether 1) textual operation information alone can predict stock returns, and 2) whether the predictive power of visual information remains robust when textual operational information is included in the specification.

Table IA.1 presents the results. Columns (2), (5), and (8) utilize only textual operational information as independent variables and show that the cumulative abnormal returns in short-run windows are not significantly correlated with both *Textual Operations Forward* and *Textual Operations Summary*, implying that the market does not significantly respond to the textual operational information embedded in slides. Additionally, in columns (3), (6), and (9), where both visual and textual operational information are included in the specification, *Textual Operations Forward* remains significantly correlated with CAR (-3, 3) and CAR (-2, 2), suggesting that the predictive power of visual information remains robust after controlling for textual operational information. Moreover, as shown in IA.2, the correlation between visual operational information and textual operational information is relatively low. Overall, this evidence suggests that visual operational information is likely incremental and valuable to market participants rather than simply redundant to the textual contents embedded in slides.

4.2 Visual Information and Future Cash Flows

Information can affect stock prices via two channels: discount rates and future cash flows (Campbell and Shiller, 1988). Given that a firm's disclosure of prospective operations is unlikely to be directly related to discount rates, we expect it to influence stock prices via the cash flow channel. We use *Earnings* (operating earnings before depreciation scaled by assets) as our primary measure for cash flows. Since revenues from new projects can be better measures of long-term profitability than short-term earnings, we also use *Sales* (scaled by total assets) as an alternative measure. Specifically, we expect a high figure for *Operations Forward* to be associated with higher future earnings and sales for the disclosing company.

We estimate the following regressions, for presentation-firm-date (i, j, t),

$$Y(t_1)_{i,j,t} = \beta_1 Operations\ Forward_{i,t} + \beta_2 Operations\ Summary_{i,t} + \gamma_1 Textual\ Controls_{i,t} + \delta Firm\ Controls_{j,t} + \alpha_{SIC2} + \mu_t + \varepsilon_{i,j,t}, \tag{7}$$

where $Y(t_1)$ is *Earnings* or *Sales* in quarter (or year) t_1 after the presentation date.

[Insert Table 5 Here]

Panel A of Table 5 reports the results of the relationship between visual operational information and future *Earnings*. While *Operations Forward* is insignificantly correlated with future *Earnings* in the 1st through 3rd quarters after the presentation (columns (1) to (3)), it significantly and positively predicts *Earnings* starting in the 4th quarter (column (4)), and continues to predict *Earnings* in the second year (column (5)). The primary reason for this may be that new operations and projects could take time to materialize and affect the firm's performance. Notably, *Operations Summary* is only significantly (and positively) related to the 1st quarter's *Earnings* following the presentation. These results are consistent with our conjecture that *Operations Forward* indeed contains longer-term information than *Operations Summary*, which also validates our machine-learning-based approach to extracting these variables from corporate presentation images.

We next examine the relation between visual operational information and revenues by replacing the outcome variable *Earnings* with *Sales*. Similar to the results for *Earnings*, Panel B of Table 5 shows that *Operations Forward* significantly and positively predicts *Sales* starting from the 4th quarter (column (4)) following the presentation, and the predictive power continues in the second year (column (5)). In addition, *Operations Summary* is not significantly correlated with long-term sales.

Overall, the evidence presented in this section indicates that forward-looking operational visual information contains important information about the disclosing firm's future revenue and profits, which in turn can influence expectations for stock prices.

4.3 Do Analysts Incorporate Visual Information?

Analysts are important information intermediaries. Investors often rely on them to provide key information through their earnings forecasts and revisions (e.g., Womack, 1996; Francis and Soffer, 1997). To the extent that *Operations Forward* contains important information regarding future earnings and sales for the disclosing company, analysts may incorporate visual information into their forecast revisions after corporate presentations. To explore this possibility, we examine the following regression at the presentation-analyst-firm-quarter level (i, k, j, t),

$$\begin{aligned} Revision\ Action_{i,k,j,t+1} = \beta_1 Operations\ Forward_{i,t} + \beta_2 Operations\ Summary_{i,t} \\ + \gamma_1 Textual\ Controls_{i,t} + \delta Firm\ Controls_{j,t} + \alpha_{SIC2} + \mu_t + \varepsilon_{i,j,t}. \end{aligned} \tag{8}$$

The dependent variable, *Revision Action*, is either *Revision Direction*, a variable equal to 1 (-1) if the analyst revises the forecast upwards (downwards) after the focal presentation date and before the subsequent earnings announcement date, or *Revision Change*, which is the change in forecasted earnings for such a revision.

[Insert Table 6 Here]

In Table 6 columns (1) and (2), we observe a significant positive relationship between *Revision Direction* and both *Operations Forward* and *Operations Summary*. Columns (3) and (4) exhibit a similar relationship between *Revision Change* and both *Operations Forward* and *Operations Summary*, which is more pronounced for *Operations Forward*. The above evidence suggests that analysts incorporate visual information into their forecasts, further validating our machine learning approach in extracting visual information from corporate presentations.

5 The AI Divide in Information Processing

In recent years, big data and AI technology have fundamentally reshaped the financial services industry, including asset management. Recent studies document how sophisticated analysts and investors are using more alternative data in recent years (e.g., Huang, Tan, and Wermers, 2020; Chi, Hwang, and Zheng, 2022). However, in light of the large volume of available data, such as may be gleaned from corporate disclosures en masse, disparities in information processing costs for different types of investors may affect their information set, trades, and market outcomes (Blankespoor et al., 2020a). Unequal access to alternative data and the capacity for utilizing such data at scale can also increase the information asymmetry between sophisticated investors and individual investors (e.g., Katona et al., 2018; Zhu, 2019).[11]

To the extent that AI talent and capability can be expensive and time-consuming to acquire (e.g., Deloitte Insights, 2020), differential access to AI may be creating an *AI divide*, which compounds information asymmetry among market participants. Executive presentations are public information that can be downloaded and read by any investor. Nonetheless, extracting visual information at a large scale would require sophisticated AI and machine learning capabilities. Still, the cost of processing visual information would be relatively low for AI-equipped financial institutions as compared to other institutions and retail investors. We conjecture that financial institutions with AI capacity are more likely to extract and utilize visual information from corporate presentations. We conduct several tests to further examine the issue.

5.1 Investor Trades, AI Capacity, and Visual Information

Our first test involves how different market participants trade a firm's stocks around corporate presentation dates. When forward-looking visual operational information is made available through executive presentations, AI-enabled institutions should trade more intensively

[11]Such a market is described as "efficiently inefficient" in Pedersen (2019).

than other investors because they have a greater capacity for processing visual information. We consider the following regression at the presentation-firm-date level (i, j, t),

$$\begin{aligned} Trades_{i,j,t} = &\ \beta_1 Operations\ Forward_i + \beta_2 Operations\ Summary_i \\ &+ \gamma Textual\ Controls_{j,t} + \delta Firm\ Controls_{i,t} \\ &+ \alpha_{SIC2} + \mu_t + \varepsilon_{i,j,t} \end{aligned} \tag{9}$$

where $Trades$ are trades made by different market participants around the presentation date t.

We first consider *AI Inst. Trade* and *Non-AI Inst. Trade* in Table 7. *AI Inst. trade* is defined as the change in the quarterly *AI Inst. Ownership* (defined in Section 3.2) during the quarter containing the presentation date.[12] Columns (1) and (2) reveal a significant and positive relationship between AI institutional trades around the presentation date and forward-looking visual operational information. In addition, AI institutional trades are insignificantly correlated with visual operational summary information. In contrast, columns (3) and (4) show that non-AI institutional trades are insignificantly correlated with both *Operations Forward* and *Operations Summary.* These results support our hypothesis that AI-equipped institutions are more likely to trade on visual operational information in presentation slides, thanks to their ability to process unstructured data at a large scale.

[Insert Table 7 Here]

In Table 8, we examine the relationship between retail trades and visual operational information. We follow Boehmer et al. (2021) and use marketable retail trade imbalance to proxy for retail trades.[13] In Panel A, the dependent variable is *Retail Order Imbalance*

[12] We proxy institutional trades by quarterly holdings change in this study due to data limitations, similar to most studies of institutional trading behavior (e.g., Griffin and Xu, 2019; Agarwal, Jiang, Tang, and Yang, 2013). The ANCerno dataset provides a subset of transaction-level institutional trades prior to around 2011, but its intersection with our sample is too small to carry out a formal study.

[13] Since we impose the availability of the retail trade measures from TAQ, the sample size is reduced compared to earlier regressions.

(Shares), the marketable retail order imbalance calculated based on the number of shares traded. Columns (1) through (4) indicate that retail trades around the presentation date are insensitive to visual operational information for the $[-3, 3]$ and $[-3, 14]$ windows. To examine whether retail investors lag in processing visual information, in columns (5) through (8), we consider retail trades for longer windows from 30 to 90 days and do not find that cumulative marketable retail order imbalance is sensitive to visual information in the long run. Panel B finds similar results for the alternative retail trade measure *Retail Order Imbalance (Trades)*, calculated based on the number of trades.

[Insert Table 8 Here]

In sum, the results above indicate that only AI-equipped institutions' trades are sensitive to the forward-looking visual operational information extracted from presentation slides, suggesting that AI-equipped institutions, as compared to non-AI-equipped institutions and retail investors, are more likely to have the ability to process unstructured image data and utilize visual information from corporate presentations.

5.2 Announcement Return and AI Institutional Ownership

If AI-equipped institutions are more able to process unstructured image data and extract visual information, we would expect to see a stronger relation between cumulative announcement returns and forward-looking visual operational information for stocks with high *AI Inst. Ownership*. To test this hypothesis, we separate stocks into two groups, based on the level of ownership by AI-equipped institutions, and we run the same specification as in Table 4 for the two groups with high and low AI institutional ownership.[14]

[Insert Table 9 Here]

[14] A stock is classified as one with high *AI Inst. Ownership* if the stock's *AI Inst. Ownership* is above the cross-sectional median for the year of presentation.

Table 9 reports the results. Columns (1) and (2) show that the cumulative announcement return in window [-3, 3] is significantly and positively correlated with forward-looking visual operational information for stocks with high *AI Inst. Ownership*. Columns (5) through (8) further show that short-term announcement returns are insignificantly correlated with *Operations Forward* for stocks with low *AI Inst. Ownership*. These findings support our hypothesis that AI-equipped institutions, as compared to non-AI-equipped institutions and retail investors, facilitate the incorporation of visual operational information into stock prices.

6 Concluding Remarks

In this paper, we use machine learning techniques to extract information from visual contents in corporate presentation slides. By using a comprehensive dataset of corporate presentation slide decks, we examine the value of visual information in predicting stock returns and study how different market participants respond to it. We find predictive power only exists in forward-looking visual operational information. Specifically, an increase in forward-looking visual operational information in presentation slides is associated with higher short-run cumulative announcement returns, and greater long-term sales and earnings following the presentation. Trades made by AI-equipped institutional holding changes are significantly and positively correlated to forward-looking visual operational information in executive presentations, while the trades of non-AI-equipped financial institutions and retail investors are not. Moreover, the relation between short-term cumulative announcement returns and forward-looking visual operational information is stronger for stocks followed by institutions with high AI utilization. The results indicate that AI-equipped institutions, as compared with other investors, are likely to have a greater capacity to process unstructured image data and extract visual information on a large scale, creating a potential *AI divide*. This can contribute to the debate about regulatory polices regarding AI. Our novel dataset and methodology can also potentially be used by researchers to study other questions in financial economics.

Appendix A: List of Variables

Variables	Definition
Operations Summary	The ratio of the number of operations-summary-related pages to the total number of pages in a presentation. Construction methodology is described in Section 2.
Operations Forward	The ratio of the number of operations-forward-related pages to the total number of pages in a presentation. Construction methodology is described in Section 2.
Textual Operations Forward	The ratio of the number of operations-forward-related pages to the total number of pages in a presentation, where operations forward pages are identified by a separate CNN + transfer model using the text in slides only.
Textual Operations Forward	The ratio of the number of operations-forward-related pages to the total number of pages in a presentation, where operations summary pages are identified by a separate CNN + transfer model using the text in slides only.
Positive Sentiment	The number of Loughran-McDonald (LM) finance-related positive words in a presentation divided by the total number of words in the presentation, expressed as a percentage.
Negative Sentiment	The number of Loughran-McDonald (LM) finance-related negative words in a presentation divided by the total number of words in the presentation, expressed as a percentage.
Uncertainty	The number of Loughran-McDonald (LM) uncertainty-related words in a presentation divided by the total number of words in the presentation, expressed as a percentage.
Constrained	The number of Loughran-McDonald (LM) constraining-related words in a presentation divided by the total number of words in the presentation, expressed as a percentage.
Forward Looking	The number of forward-looking words in a presentation divided by the total number of words in the presentation, expressed as a percentage. We construct a comprehensive set of forward-looking words including all dictionaries in Li (2010), Muslu et al.(2015), Bozanic et al.(2018), and Grewal et al.(2019)
Text Length	The total number of words in the presentation.
Size	The natural logarithm of the market capitalization.
Book-to-Market	Total assets over market cap plus the book value of liabilities.
Turnover	The monthly average of the ratio of trading volume to shares outstanding, multiplied by 12.
NASDAQ	NASDAQ indicator that equals one if the stock is traded on the NASDAQ.
Inst. Ownership	The ratio of the total shares of institutional ownership to shares outstanding.
Prior Return	Average FF3 alpha in window [-71,-11] in terms of trading days prior to presentations.

(continued)

Variables	Definition
Revision Direction	A dummy variable that equals one if the analyst changes their forecast for the focal stock during the period between the presentation date and the subsequent earnings announcement date.
Revision Change	The magnitude of the change an analyst makes to their forecast for the focal stock during the period between presentation date and the subsequent earnings announcement date.
General Experience	An analyst's general experience.
Industry Experience	An analyst's industry experience.
Star Analyst	A dummy equal to one if the analyst is a star analyst
Brokerage Size	The size of the brokerage firm that the analyst belongs to.
Earnings	Operating earnings before depreciation scaled by assets.
Sales	Sales scaled by total assets).
AI Inst. Ownership	The AI institutional ownership for a given stock. Please refer to Section 2 for detailed construction method.
AI Inst. Trades	The holdings change among AI-equipped financial institutions. Please refer to Section 2 for detailed construction method.
Aggregate Inst. Size	The sum of the sizes of all institutions holding the focal stock.
# of Institutions	The number of institutions holding the focal stock.
Retail Order Imbalance (Shares)	Marketable retail order imbalance trade volume. Please refer to Section 2 for detailed construction method.
Retail Order Imbalance (Trades)	Marketable retail order imbalance trades. Please refer to Section 2 for detailed construction method.

Appendix B: Convolutional Neural Networks

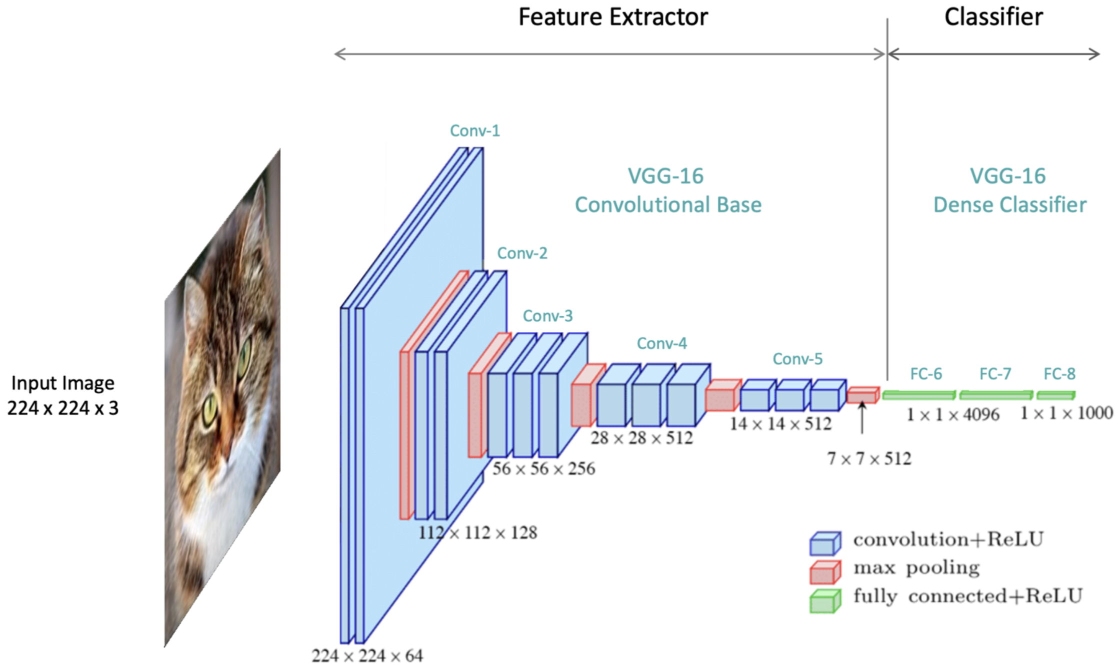

Figure 1: Illustration of the VGG CNN Model. (Source: Learnopencv.com)

In this section, we review details of the Convolutional Neural Network (CNN or ConvNet) and transfer learning models. One limitation of neural networks is that they do not explicitly consider the grid topology or variable dependency, and so fail to model complex interactions between predictors in our context. To overcome this challenge, we apply the CNNs as a feature extractor. CNNs have been tremendously successful in applications ranging from image to text classifications. The CNN architecture we used consists of an *input layer* that takes the pixels of the image into its nodes, one or more *convolutional layers* that utilize the convolution operator to model the complicated nonlinear interactions between the input pixels and generate intermediate features that capture progressively higher-level information from the image (such as shapes and objects), one or more *batch normalization layers* and *max-pooling layers* that reduce the computation time, extract the local dependency across predictors and yield a robust model, and a final *fully connected layer* and *output layer* to finish the classification and makes prediction for the category the image belongs to. Figure 1 provides an illustration of the VGG-16 CNN network, on which our final model is based.

Convolutional Layer. Based on the input image block $A_{i,j}$, the convolutional layer applies the convolution operator to generate the intermediate features as follows:

$$c_{i,j} = \sigma(W^{conv} * A_{i,j} + b^{conv}), \tag{A1}$$

where $\sigma(\cdot)$ is the ReLU activation function, W^{conv} and b^{conv} respectively represent the weight of the filter and the bias, and $*$ is the convolution operator.

Max-Pooling Layer. The pooling operation replaces the output $c_{i,j}$ of the convolutional layer at a certain location with a summary statistic of the nearby neighborhood. The max-pooling operation reports the maximum output within a rectangular neighborhood,

$$m_{i,j} = \max(c_{k,l}), \quad k = i-1, i, i+1; \quad l = j-1, j, j+1; \tag{A2}$$

Fully Connected Layer. The output $M = [m_{i,j}]$ from the max-pooling layer is flattened and concatenated with a vector of text input into a vector F, which is further transformed in the fully connected (FC) layer,

$$fc^k = \sigma(W_k^{fc} F + b_k^{fc}), \quad k = 1, \cdots, K \tag{A3}$$

where $\sigma(\cdot)$, W_k^{fc}, b_k^{fc}, and K represents the ReLU function, input weights, biases and the number of classification categories. The nonlinearity introduced by the fully connected layer further enhances the model's ability to learn the internal nonlinearity and to generate better prediction results.

Output Layer. The final output layer consists of K neurons, which represent the probability of each category. Based on the output of the fully connected layer $fc = [fc^1, \cdots, fc^K]$, the softmax activation function yields

$$p(k) = Softmax(W_k fc + b_k), \quad k = 1, \cdots, K \tag{A4}$$

Transfer CNN. The weights of the CNN network need to be estimated using the training data. When the number of observations is large enough, the neural network can learn and produce a model with high accuracy. When the data size is not sufficiently large, it is possible for the model to "learn" from similar image classification problems. We refer to the CNN model that was learned by combining the knowledge learned from our data and other existing data sets as the *Transfer CNN* model. The basic premise of transfer learning is simple: take a model trained on a large data set and transfer its knowledge to a smaller data set. For our specific training task, we build on the pre-trained VGG16 (Simonyan and Zisserman, 2014) using ImageNet data (Deng et al., 2009), by keeping the network structure and the trained parameters from all the layers except the fully connected layer. That is, we train a VGG16 network by treating only the parameters in the fully connected layer as unknown.

Appendix C: Examples of Visual Information

Operations Forward

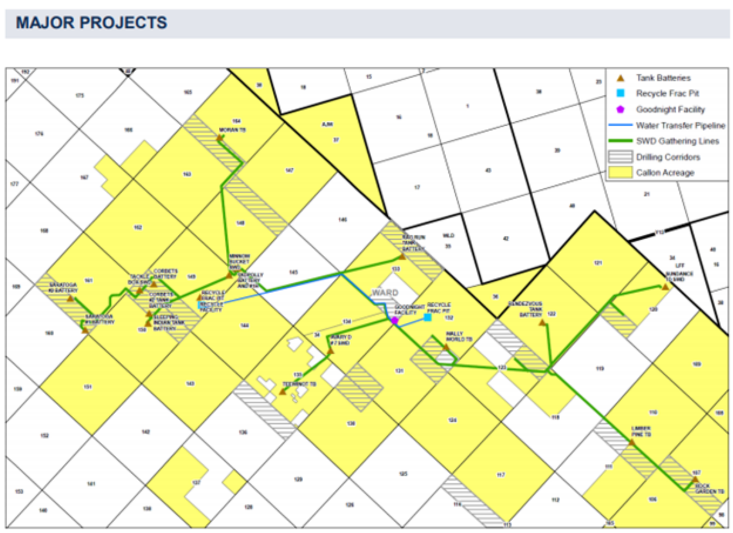

Operations Summary

- Taipei's 101 Mall

Others

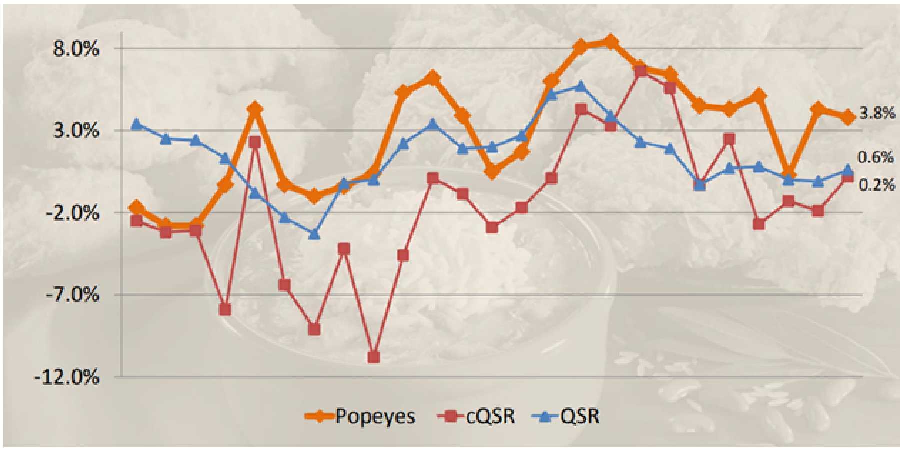
8.0%
3.0%
-2.0%
-7.0%
-12.0%
3.8%
0.6%
0.2%
Popeyes
cQSR
QSR

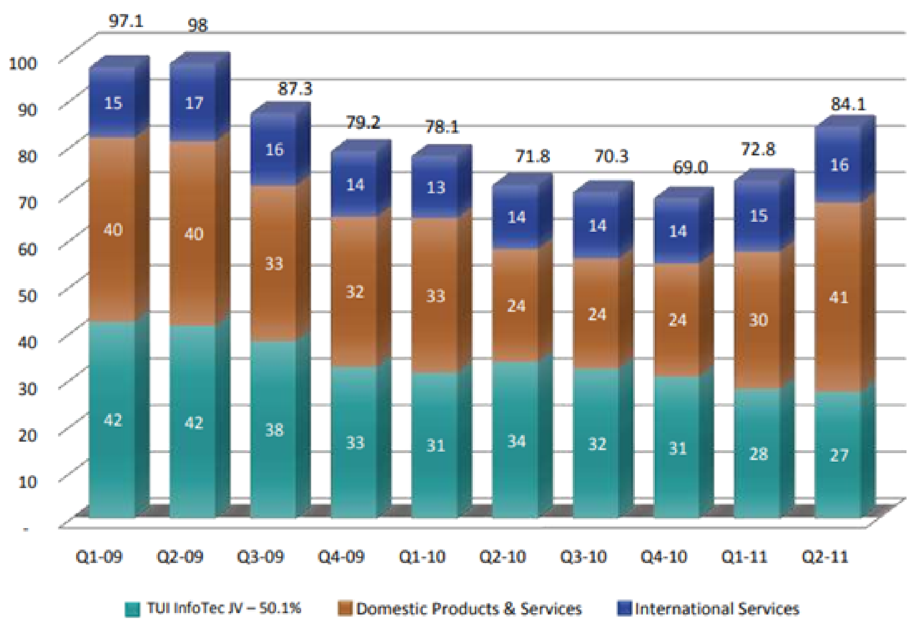
100
90
80
70
60
50
40
30
20
10
-
97.1
98
87.3
79.2
78.1
71.8
70.3
69.0
72.8
84.1
15 17 16 14 13 14 14 14 15 16
40 40 33 32 33 24 24 24 30 41
42 42 38 33 31 34 32 31 28 27
Q1-09 Q2-09 Q3-09 Q4-09 Q1-10 Q2-10 Q3-10 Q4-10 Q1-11 Q2-11
TUI InfoTec JV – 50.1%
Domestic Products & Services
International Services

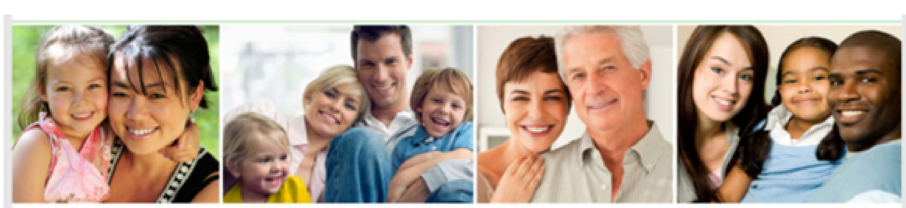

Appendix D: AI Labor Stock Estimation

Step 1: Identify AI-related Skills

We first identify AI-related skills following Babina et al. (2021). Intuitively, jobs requiring a certain ability related to AI should also require other core AI skills, and a skill that frequently co-occurs with unambiguous core AI skills might be related to AI. The calculation method is summarized as follows. We start by creating a sample of job postings that require four core AI skills—"Machine Learning," "Natural Language Processing," "Computer Vision," and "Artificial Intelligence." We then measure the AI-relatedness of any given skill as the fraction of job postings requiring that skill, as well as four AI skills mentioned to the total number of job postings requiring that skill. Specifically, for each skill s, the AI-relatedness w of that skill is:

$$w_s^{AI} = \frac{\#\ of\ jobs\ with\ skill\ s\ and\ \{AI, ML, NLP, or\ CV\}\ in\ job\ title\ or\ in\ skills}{\#\ of\ jobs\ with\ skill\ s} \tag{A5}$$

Step 2: Identify AI-related Job Postings

After identifying AI-related skills, we calculate AI-relatedness for each job posting by taking the average of the AI-relatedness of the skills associated with that job posting. This gives us a continuous AI-relatedness measure (ranging from 0 to 1) for each job posting. Specifically, letting N denote the number of required skills listed for job posting j, our job-level AI-relatedness measure is:

$$w_j^{AI} = \frac{1}{N}\sum_{s=1}^{N} w_s^{AI}. \tag{A6}$$

Step 3: AI-relatedness for Institutions

After obtaining AI-relatedness for each job posting, we calculate AI-relatedness at the institution-year level. To be specific, for institution i in year t, its AI-relatedness is the number of AI-related job postings posted by the institution scaled by the total number of job postings:

$$p_{i,t}^{AI} = \frac{1}{K}\sum_{j=1}^{K} 1_j(w_j^{AI} > \varphi), \tag{A7}$$

where φ is the cutoff for whether a job is AI job.[15] We conduct fuzzy name matching complemented with manual verification to match Burning Glass institutions to Thomson Reuters 13F institutions.[16]

Step 4: Accumulated AI Labor Stock

We estimate accumulated AI labor stock following a similar method to Abis and Veldkamp (2020). First, we use data from the Bureau of Labor Statistics to estimate the probability that a vacancy is filled and the probability that an employed worker separates from their job.[17] Then we calculate the labor stock as follows, for institution-year (i, t),

$$l_{i,t}^{AI} = l_{i,t-1}^{AI}\left(1 - s_t^{AI}\right) + p_{i,t}^{AI} \times h_t^{AI}, \tag{A8}$$

where $l_{i,t}^{AI}$ is the AI labor stock, s_t^{AI} and h_t^{AI} are the separation rate and vacancy fill rate for the financial services sector[18], and $p_{i,t}^{AI}$ is the AI-relatedness calculated in Step 3. For example, if fund A has 20 AI employees in year 2016, and it posts 10 AI job postings in year 2017, with the estimated average separation rate of 0.1 and hiring rate of 0.6, then fund A's AI labor stock in 2017 is calculated as $20 \times (1 - 0.1) + 10 \times 0.6 = 24$.

[15]Intuitively, a higher cutoff imposes stricter criteria for AI jobs. In our main analysis, we set φ to 0.1.

[16]We keep asset management companies with type code = 3, 4, 5 in TR 13F data for our main analysis. Approximately 30% of TR 13F institutions can be matched to Burning Glass in our sample.

[17]Data are available at `https://www.bls.gov/jlt/#data`.

[18]Finance and Insurance (NAICS 52) industry according to the BLS classification.

References

Abis, S., Veldkamp, L. (2020). The changing economics of knowledge production. Available at SSRN 3570130.

Abis, S. (2020). Man vs. machine: Quantitative and discretionary equity management. Machine: Quantitative and Discretionary Equity Management (October 23, 2020).

Acemoglu, D., Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488-1542.

Acemoglu, D., Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188-2244.

Acemoglu, D., Autor, D., Hazell, J., Restrepo, P. (2020). AI and jobs: Evidence from online vacancies (No. w28257). National Bureau of Economic Research.

Agrawal, A., Gans, J. S., Goldfarb, A. (2019). Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2), 31-50.

Agarwal, V., Jiang, W., Tang, Y., Yang, B. (2013). Uncovering hedge fund skill from the portfolio holdings they hide. The Journal of Finance, 68(2), 739-783.

Aghion, P., Jones, B. F., Jones, C. I. (2018). Artificial intelligence and economic growth. In The economics of artificial intelligence: An agenda (pp. 237-282). University of Chicago Press.

Babina, T., Fedyk, A., He, A. X., Hodson, J. (2021). Artificial Intelligence, Firm Growth, and Product Innovation. Firm Growth, and Product Innovation (November 9, 2021).

Ben-Rephael, A., Ronen, J., Ronen, T., Zhou, M. (2021). Do images provide relevant information to investors? An exploratory study. An Exploratory Study (November 18, 2021).

van Binsbergen, Jules H., Xiao Han, and Alejandro Lopez-Lira, 2022, Man vs. machine learning: The term structure of earnings expectations and conditional biases, *Review of Financial Studies* forthcoming.

Blankespoor, E., Hendricks, B. E., Miller, G. S. (2017). Perceptions and price: Evidence from CEO presentations at IPO roadshows. *Journal of Accounting Research*, 55(2), 275-327..

Blankespoor, E., deHaan, E., Marinovic, I. (2020). Disclosure processing costs, investors' information choice, and equity market outcomes: A review. *Journal of Accounting and Economics*, 70(2-3), 101344.

Blankespoor, E., Hendricks, B. E., Miller, G. S. (2020). The pitch: Managers' disclosure choice during IPO roadshows. *The Accounting Review*.

Boehmer, E., Jones, C. M., Zhang, X., Zhang, X. (2021). Tracking retail investor activity. *The Journal of Finance*, 76(5), 2249-2305.

Bozanic, Z., Roulstone, D. T., Van Buskirk, A. (2018). Management earnings forecasts and other forward-looking statements. *Journal of Accounting and Economics*, 65(1), 1-20.

Bradley, D., Jame, R., Williams, J. (2021). Non-Deal Roadshows, Informed Trading, and Analyst Bias. Journal of Finance. *The Journal of Finance.*

Brogaard, Jonathan and Abalfazl Zareei, 2022, Machine learning and the stock market, *Journal of Financial and Quantitative Analysis*, forthcoming.

Brynjolfsson, E., Mitchell, T., Rock, D. (2018, May). What can machines learn, and what does it mean for occupations and the economy?. In AEA papers and proceedings (Vol. 108, pp. 43-47).

Brynjolfsson, E., Rock, D., Syverson, C. (2018). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In The economics of artificial intelligence: An agenda (pp. 23-57). University of Chicago Press.

Bushee, B.J., M. J. Jung, and G. S. Miller, 2011, Conference Presentations and the Disclosure Milieu, *Journal of Accounting Research* 49 (5), 1163-1192.

Campbell, J. Y., Shiller, R. J. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies*, 1(3), 195-228.

Cao, S., Cong, L. W., Yang, B. (2019). Financial reporting and blockchains: Audit pricing, misstatements, and regulation. Misstatements, and Regulation (June 2019).

Cao, S., Jiang, W., Wang, J. L., Yang, B. (2021). From man vs. machine to man+ machine: The art and ai of stock analyses (No. w28800). National Bureau of Economic Research.

Cao, S., Jiang, W., Yang, B., Zhang, A. L. (2020). How to talk when a machine is listening: Corporate disclosure in the age of AI (No. w27950). National Bureau of Economic Research.

Cao, Kai, and Haifeng You, 2021, Fundamental analysis via machine learning, Working paper.

Cong, Lin William, Tengyuan Liang, Baozhong Yang, and Xiao Zhang, 2020, Analyzing textual information at scale, in Kashi R. Balachandran, ed.: *Information to Facilitate Efficient Decision Making: Big Data, Blockchain and Relevance* (World Scientific Publishers, Singapore).

Cong, Lin William, Ke Tang, Jingyuan Wang, and Yang Zhang, 2022, AlphaPortfolio: Direct construction through reinforcement learning and interpretable AI, Working paper.

Chen, J., Hong, H., Stein, J. C. (2002). Breadth of ownership and stock returns. *Journal of financial Economics*, 66(2-3), 171-205.

Chen, Luyang, Markus Pelger, and Jason Zhu, 2022, Deep learning in asset pricing, *Management Science*, forthcoming.

Chi, F., Hwang, B. H., Zheng, Y. (2022). The Use and Usefulness of Big Data in Capital Markets. Nanyang Business School Research Paper, (22-01).

Christensen, T. E., Fronk, K., Lee, J. A., Nelson, K. K. (2021). Data Visualization in 10-K Filings. Available at SSRN 3748711.

Coleman, B., Merkley, K. J., Pacelli, J. (2021). Human versus machine: A comparison of robo-analyst and traditional research analyst investment recommendations. The Accounting Review, Forthcoming.

Cong, L. W., Tang, K., Wang, J., Zhang, Y. (2021). AlphaPortfolio: Direct construction through deep reinforcement learning and interpretable AI. SSRN Electronic Journal. https://doi.org/10.2139/ssrn, 3554486.

Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., Fei-Fei, L. (2009, June). Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition (pp. 248-255). Ieee.

Deng, W., Gao, L., Hu, B., Zhou, G. (2020). Seeing is Believing: Annual Report 'Graphicity' and Stock Returns Predictability. Available at SSRN 3723126.

Ellis, J. A., Gerken, W. C., Jame, R. (2021). The Meeting before the Meeting: Governance Roadshows and Mutual Fund Voting. Available at SSRN..

Francis, J., Soffer, L.(1997). The relative informativeness of analysts' stock recommendations and earnings forecast revisions. *Journal of Accounting Research*, 35(2), 193-211.

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *The Journal of Finance*, 77(1), 5-47.

Goodfellow, I., Y., Bengio, and A., Courville. 2016. *Deep Learning*, MIT Press.

Green, T. C., Jame, R., Markov, S., Subasi, M. (2014). Broker-hosted investor conferences. *Journal of Accounting and Economics*, 58(1), 142-166.

Grennan, J., Michaely, R. (2020). Artificial intelligence and high-skilled work: Evidence from analysts. Swiss Finance Institute Research Paper, (20-84).

Grennan, J., Michaely, R. (2021). Fintechs and the market for financial analysis. *Journal of Financial and Quantitative Analysis*, 56(6), 1877-1907.

Grewal, R., Lilien, G. L., Petersen, J. A., Wuyts, S. (2022). Business-to-business marketing: Looking back, looking forward. In Handbook of business-to-business marketing (pp. 2-11). Edward Elgar Publishing.

Griffin, J. M., Xu, J. (2009). How smart are the smart guys? A unique view from hedge fund stock holdings. *The Review of Financial Studies*, 22(7), 2531-2570.

Gu, S., Kelly, B., Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223-2273.

Hanley, Kathleen Weiss, and Gerard Hoberg, 2010, The information content of IPO prospectuses, *Review of Financial Studies* 23, 2821–2864.

Hu, A., Ma, S. (2021). Persuading investors: A video-based study (No. w29048). National Bureau of Economic Research.

Hupfer, S. (2020). Talent and Workforce Effects in the Age of AI. *Deloitte Insights*.

Hoberg, G., and G. Phillips, 2016. Text-Based Network Industries and Endogenous Product Differentiation, *Journal of Political Economy* 124(5), 1423-1465.

Huang, Alan Guoming, Hongping Tan, Russ Wermers, 2020, Institutional trading around corporate news: Evidence from textual analysis, *Review of Financial Studies*, Forthcoming.

An introduction to statistical learning (Vol. 112, p. 18). New York: springer.

Jegadeesh, N., and D. Wu, 2013. Word Power: A New Approach for Content Analysis, *Journal of Financial Economics* 110, 712-729.

Jiang, F., Lee, J., Martin, X. and Zhou, G., 2019. Manager sentiment and stock returns. *Journal of Financial Economics*, 132(1), 126-149.

Katona, Z., Painter, M., Patatoukas, P. N., Zeng, J. (2018, July). On the capital market consequences of alternative data: Evidence from outer space. In 9th Miami Behavioral Finance Conference.

Lee, C. M., Ready, M. J. (1991). Inferring trade direction from intraday data. *The Journal of Finance*, 46(2), 733-746.

Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., Fei-Fei, L. (2009, June). Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition (pp. 248-255). Ieee.

Li, F. 2010. The Information Content of Forward-Looking Statements in Corporate Filings-A Naive Bayesian Machine Learning Approach. *Journal of Accounting Research* 48 (5), 1049-1102.

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

Loughran, Tim, and Bill McDonald, 2016, Textual analysis in accounting and finance: A survey, *Journal of Accounting Research* 54, 1187–1230.

Mayew, W. J., Venkatachalam, M. (2012). The power of voice: Managerial affective states and future firm performance. *The Journal of Finance*, 67(1), 1-43.

Muslu, V., Radhakrishnan, S., Subramanyam, K. R., Lim, D. (2015). Forward-looking MDA disclosures and the information environment.*Management Science*, 61(5), 931-948.

Obaid, K., Pukthuanthong, K. (2022). A picture is worth a thousand words: Measuring investor sentiment by combining machine learning and photos from news. Journal of Financial Economics, 144(1), 273-297.

Pedersen, L. H. (2019). Efficiently inefficient: how smart money invests and market prices are determined. Princeton University Press.

Pratt, L. Y., 1993. Discriminability-based transfer between neural networks, NIPS Conference: Advances in Neural *Information Processing Systems* 5, Morgan Kaufmann Publishers, 204â211.

Rajat, R, A.Y., Ng, and D., Koller, 2006. Constructing Informative Priors using Transfer Learning, *Twenty-third International Conference on Machine Learning.*

de Silva, Tim, and David Thesmar, 2021, Noise in expectations: Evidence from analyst forecasts, NBER Working paper 28963.

Simonyan, K., Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.

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

Womack, K. L. (1996). Do brokerage analysts' recommendations have investment value?. *The journal of finance*, 51(1), 137-167.

Zhu, C. (2019). Big data as a governance mechanism. *The Review of Financial Studies*, 32(5), 2021-2061.

Figure 2: Time Series of Corporate Presentations

This figure plots the annual number of presentations (as the bar plots and corresponding to the left axis) and the ratio of the number of presentations with *Operations Forward* images to the number of all presentations (as the line plot and corresponding to the right axis) in our sample from 2006 to 2018. A presentation contains *Operations Forward* images if any slide in the presentation is classified to the *Operations Forward* category.

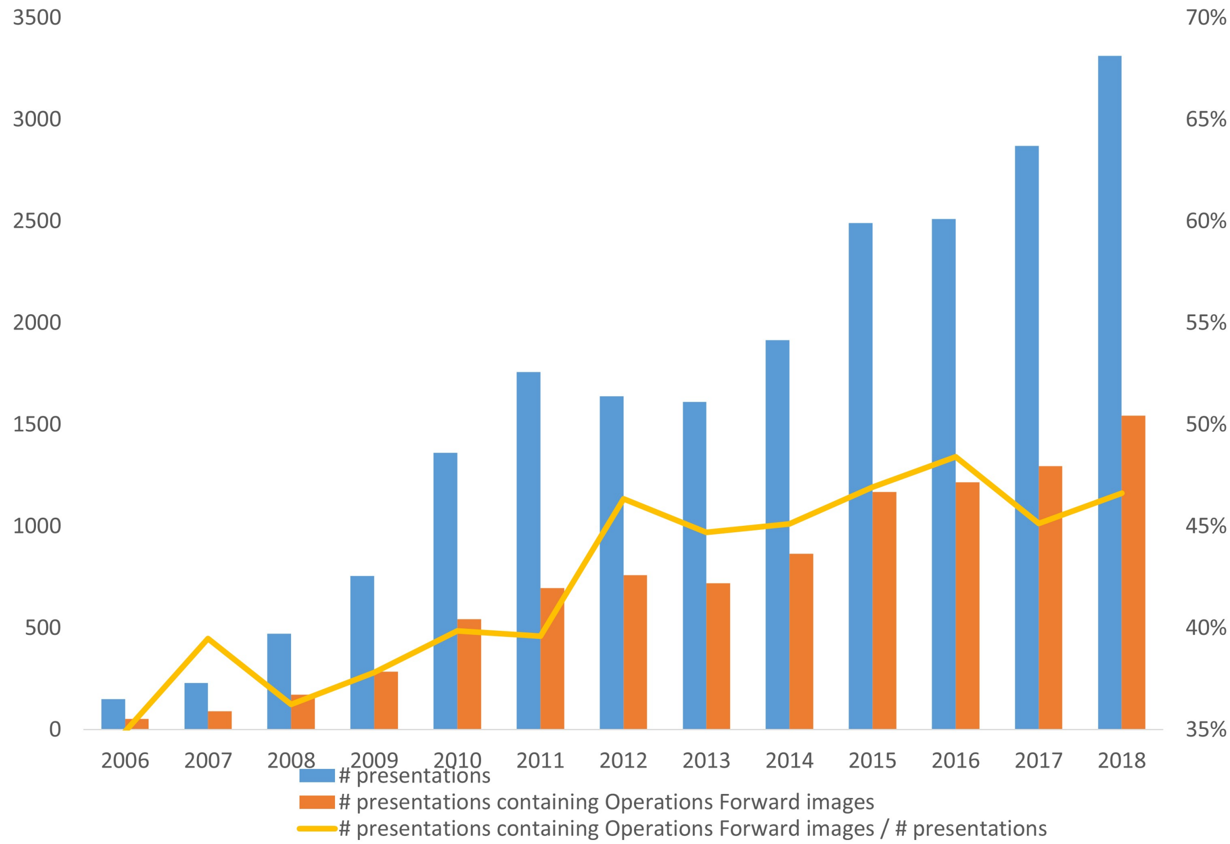

Figure 3: Monthly Distribution of Corporate Presentations

This figure plots the distributions of presentation dates and 10-K/10-Q filing dates in the calendar months.

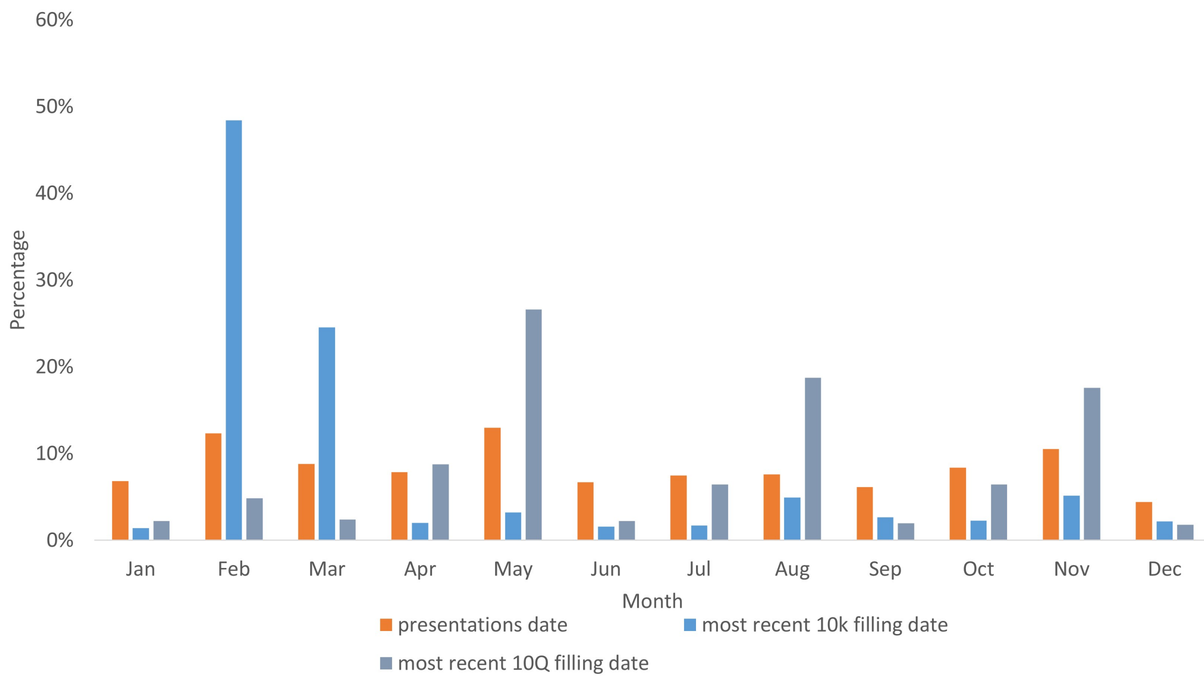

Table 1: Performance of Machine Learning Models

This table summarizes the performance of different machine learning methods to identify and classify corporate image information. We report the out-of-sample performance calculated by the ten-fold cross-validation method. *Accuracy* is one minus the ratio of incorrect category predictions to total observations. For each category, *Precision* is the ratio of true positives to the sum of true positives and false positives; *Recall* is the ratio of true positives to the sum of true positives and false negatives. We calculate *Precision* and *Recall* for each category and then average across the three categories. *F1 Score* is the harmonic mean of *Precision* and *Recall*. Model parameters are selected to maximize the *F1 score* in cross-validation.

Model	*CNN* (%)	*CNN + Text* (%)	*Transfer CNN* (%)	*Transfer CNN + Text* (%)
Accuracy	75.0	76.5	77.0	80.0
Precision	77.3	70.7	77.5	78.9
Recall	75.0	76.5	77.0	80.0
F1 Score	76.1	73.5	77.2	79.4

Table 2: Summary Statistics

This table presents the descriptive statistics of key variables used in the paper. Panel A reports the statistics for presentation-level variables (i.e., variables constructed using information in slides); Panel B reports the statistics for firm-level variables. All variables are defined as in Appendix A. *AI Inst. Ownership* is constructed as in Section 2.

Variables	Mean	Median	Std	P25	P75	N
Panel A: Presentation-level Variables						
Operations Summary	0.110	0.075	0.122	0.000	0.167	17,277
Operations Forward	0.036	0.000	0.058	0.000	0.056	17,277
Positive Sentiment	0.014	0.012	0.008	0.009	0.017	17,277
Negative Sentiment	0.012	0.011	0.005	0.008	0.014	17,277
Uncertainty	0.010	0.009	0.006	0.006	0.013	17,277
Constrained	0.003	0.003	0.002	0.001	0.004	17,277
Forward Looking	0.043	0.040	0.019	0.030	0.053	17,277
Text Length	3088	2480	2381	1586	3864	17,277
Panel B: Firm-level Variables						
Size	13.10	3.12	30.33	1.08	10.48	17,277
Book-to-market	0.615	0.529	0.421	0.306	0.833	17,277
Prior Return	0.000	0.000	0.001	-0.000	0.001	17,277
Nasdaq	0.311	0.000	0.463	0.000	1.000	17,277
Turnover	1557	1268	1213	8721	1888	17,277
Inst. Ownership	0.775	0.819	0.200	0.705	0.914	17,277
AI Inst. Ownership	0.377	0.380	0.134	0.298	0.457	17,277

Table 3: Top Industries by Operations Forward

This table lists the top ten 2-digit SIC industries with forward-looking operational information. In Panel A, industries are ranked by the number of presentations with at least one *Operations Forward* slide, scaled by the number of all types of presentations. In Panel B, industries are ranked by the industry average of the variable *Operations Forward*, which is the number of slides in this category relative to the total number of slides. We exclude industries with fewer than 50 presentations in the sample.

Panel A: Top 10 industries by proportion of *Operations Forward* presentations

SIC2	Name	% of *Operations Forward* Presentations
10	Metal, Mining	78%
13	Oil & Gas Extraction	78%
45	Transportation by Air	72%
51	Wholesale Trade (non-durable Goods)	71%
31	Leather & Leather Products	59%
79	Amusement & Recreation Services	58%
49	Electric, Gas, & Sanitary Services	57%
23	Apparel & Other Textile Products	55%
55	Automative Dealers & Service Stations	53%
20	Food & Kindred Products	53%

Panel B: Top 10 industries by average proportion of *Operations Forward* slides

SIC2	Name	Average *Operations Forward*
10	Metal, Mining	8.7%
40	Railroad Transportation	8.3%
13	Oil & Gas Extraction	7.8%
44	Water Transportation	7.7%
16	Heavy Construction, Except Building	6.9%
45	Transportation by Air	6.5%
51	Wholesale Trade (non-durable Goods)	5.9%
32	Stone, Clay, & Glass Products	5.6%
33	Primary Metal Industries	5.2%
20	Food & Kindred Products	5.0%

Table 4: Visual Information and Market Response

This table examines the relation between visual information in corporate presentation slides and announcement returns. The dependent variables are *Cumulative Announcement Returns* in different windows, where returns are adjusted by *Market*, *Size*, and *Book-to-Market*. The sample period is from 2006 to 2018. *Operations Forward* and *Operations Summary* measure the forward-looking operational visual information and backward-looking operational visual information in presentation slides. All variables are defined as in Appendix A. Industries are defined by 2-digit SIC codes. Heteroskedasticity-consistent standard errors (in parentheses) are clustered at the firm level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels (two-tailed), respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variables	*CAR(-3,3)*		*CAR(-2,2)*		*CAR(4,13)*	
Operations Forward	0.027**	0.024*	0.020	0.021*	-0.016	-0.016
	(0.014)	(0.014)	(0.013)	(0.013)	(0.011)	(0.011)
Operations Summary	-0.009	-0.009	-0.003	-0.002	0.005	0.005
	(0.008)	(0.008)	(0.006)	(0.007)	(0.006)	(0.006)
Controls Variables						
Positive Sentiment	0.685***	0.613***	0.649***	0.580***	-0.280	-0.280
	(0.145)	(0.154)	(0.121)	(0.129)	(0.177)	(0.177)
Negative Sentiment	-0.449**	-0.300	-0.391**	-0.277	0.286*	0.286*
	(0.204)	(0.203)	(0.194)	(0.193)	(0.161)	(0.161)
Uncertainty	0.179	0.028	-0.057	-0.167	0.289**	0.289**
	(0.158)	(0.159)	(0.143)	(0.146)	(0.118)	(0.118)
Constrained	0.139	-0.054	0.085	-0.004	-0.134	-0.134
	(0.336)	(0.355)	(0.305)	(0.323)	(0.296)	(0.296)
Forward-Looking	-0.154**	-0.126**	-0.164***	-0.143**	-0.024	-0.024
	(0.060)	(0.063)	(0.056)	(0.058)	(0.047)	(0.047)
Text Length	0.000	0.000	-0.001	0.000	0.001	0.001
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)

(continued)

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variables	*CAR(-3,3)*		*CAR(-2,2)*		*CAR(4,13)*	
Size	-0.003***	-0.003***	-0.002***	-0.002***	-0.002***	-0.002***
	(0.001)	(0.001)	(0.000)	(0.001)	(0.000)	(0.000)
Book-to-Market	0.001	0.001	0.001	0.002	-0.003***	-0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Turnover	0.000	0.003	-0.001	0.002	-0.004***	-0.004***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
Nasdaq	-0.001	-0.002	0.002	-0.001	-0.003*	-0.003*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
InstOwnership	0.005	-0.002	0.011**	0.003	-0.002	-0.002
	(0.006)	(0.006)	(0.005)	(0.006)	(0.004)	(0.004)
Prior Return	-1.328	-1.531	-2.053**	-2.106**	-0.475	-0.475
	(1.199)	(1.188)	(1.029)	(1.027)	(0.819)	(0.819)
Industry FE	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,326	17,277	17,326	17,277	17,275	17,275
R-squared	0.009	0.016	0.009	0.016	0.017	0.019

Table 5: Visual Information and Future Cash Flows

This table reports the relation between visual information in corporate presentation slides and firms' future cash flows. Panels A and B examine whether visual information can predict *Earnings* and *Sales*, respectively. *Operations Forward* and *Operations Summary* measure the forward-looking operational visual information and backward-looking operational visual information in presentation slides. We include the same set of control variables as in Table 4. All variables are defined as in Appendix A. Heteroskedasticity-consistent standard errors (in parentheses) are clustered at the firm level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels (two-tailed), respectively.

Panel A: Visual Information and Earnings

	(1) 1st Qtr	(2) 2nd Qtr	(3) 3rd Qtr	(4) 4th Qtr	(5) 2nd Year
Operations Forward	-0.002	0.002	0.003	0.008**	0.032*
	(0.004)	(0.005)	(0.004)	(0.004)	(0.017)
Operations Summary	0.006***	0.002	0.005	-0.002	0.002
	(0.002)	(0.002)	(0.003)	(0.002)	(0.008)
Controls	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	16,884	16,808	16,722	16,134	13,129
Adjusted R-squared	0.521	0.476	0.602	0.626	0.715

Panel B: Visual Information and Sales

	(1) 1st Qtr	(2) 2nd Qtr	(3) 3rd Qtr	(4) 4th Qtr	(5) 2nd Year
Operations Forward	0.003	0.009	0.008	0.022***	0.125***
	(0.007)	(0.01)	(0.008)	(0.008)	(0.042)
Operations Summary	0.009**	0.000	0.008	-0.001	0.025
	(0.004)	(0.005)	(0.005)	(0.004)	(0.021)
Controls	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	17,230	17,192	17,131	16,539	13,562
Adjusted R-squared	0.965	0.952	0.959	0.967	0.955

Table 6: Visual Information and Analyst Revision

This table presents regression results for the relation between analyst revision and visual information in corporate presentation slides. *Revision Direction* is equal to 1 (-1) if the analyst revises up (down) the forecast after the focal presentation date and before the subsequent earning announcement date, and *Revision Change* is the change in forecasted earnings for such a revision. The independent variables of main interest are *Operations Forward* and *Operations Summary*, which measure the forward-looking operational visual information and backward-looking operational visual information in presentation slides. We include the same set of control variables as in Table 4. All variables are defined as in Appendix A. Heteroskedasticity-consistent standard errors (in parentheses) are clustered at the firm level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels (two-tailed), respectively.

	(1)	(2)	(3)	(4)
Dependent Variables	*Revision Direction*		*Revision Change*	
Operations Forward	0.419***	0.655***	0.393***	0.366***
	(0.085)	(0.081)	(0.121)	(0.117)
Operations Summary	0.211***	0.281***	0.104*	0.106*
	(0.045)	(0.043)	(0.062)	(0.060)
Controls Variables				
# of Prior Forecasts	-2.343***	-2.588***	-4.654***	-4.664***
	(0.160)	(0.121)	(0.225)	(0.168)
General Experience	0.001**	0.001	0.000	0.000
	(0.000)	(0.000)	(0.001)	(0.001)
Industry Experience	-0.003***	-0.002***	0.000	0.000
	(0.000)	(0.000)	(0.001)	(0.001)
Firm Experience	-0.007***	-0.007***	-0.006***	-0.006***
	(0.000)	(0.000)	(0.000)	(0.000)
Star Analyst	0.054	0.034	-0.001	-0.023
	(0.070)	(0.070)	(0.102)	(0.106)
Brokerage Size	-0.001***	-0.001***	-0.001***	-0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Controls	Yes	Yes	Yes	Yes
Industry FE	Yes	No	Yes	No
Year FE	Yes	Yes	Yes	Yes
Observations	337,953	337,953	337,944	337,944
Adjusted R-squared	0.411	0.409	0.643	0.640

Table 7: Visual Information and Institutional Trades

This table examines the relation between institutional holding change and visual information in corporate presentation slides. *AI Inst. Trade* and *Non-AI Inst. Trade* are changes in aggregate institutional holdings for AI-equipped and non-AI-equipped financial institutions, defined in Section 3.2. We restrict our sample to events from 2010 to 2018 due to limited data coverage in Burning Glass before 2010. The independent variables of main interest are *Operations Forward* and *Operations Summary*, which measure forward-looking operational visual information and backward-looking operational visual information in presentation slides. We include the same set of control variables as in Table 4. We also control for *Aggregate Institutional Ownership* of the stock and *# of Institutions* holding the stock. All variables are defined in Appendix A. Heteroskedasticity-consistent standard errors (in parentheses) are clustered at the firm level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels (two-tailed), respectively.

	(1)	(2)	(3)	(4)
Dependent Variables	*AI Inst. Trade*		*Non-AI Inst. Trade*	
Operations Forward	0.011**	0.012**	-0.001	-0.001
	(0.005)	(0.006)	(0.004)	(0.004)
Operations Summary	0.003	-0.001	0.002	0.002
	(0.003)	(0.003)	(0.003)	(0.003)
Controls	No	Yes	No	Yes
Industry FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	15,850	15,850	15,850	15,850
Adjusted R-squared	0.001	0.003	0.001	0.003

Table 8: Visual Information and Marketable Retail Order Imbalance

This table reports the results on the relation between marketable retail order imbalance and visual information in corporate presentation slides. In Panel A, the dependent variable is the marketable retail order imbalance calculated based on the number of shares traded; in Panel B, the dependent variable is the marketable retail order imbalance calculated based on the number of trades. Detailed definitions are given in Section 2. *Operations Forward* and *Operations Summary* measure the forward-looking operational visual information and backward-looking operational visual information in presentation slides, respectively. We include the same set of control variables as in Table 4. All variables are defined in Appendix A. Heteroskedasticity-consistent standard errors (in parentheses) are clustered at the firm level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels (two-tailed), respectively.

Panel A: Marketable Retail Trade Imbalance based on # of Shares Traded

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable	*Retail Order Imbalance (shares)*							
Window	[-3,3]		[-3,14]		[-3,30]		[-3,90]	
Operations Forward	0.028	0.029	0.022	0.023	0.039	0.040	0.033	0.033
	(0.043)	(0.043)	(0.036)	(0.036)	(0.027)	(0.027)	(0.023)	(0.023)
Operations Summary	-0.028	-0.029	-0.009	-0.010	0.006	0.006	0.004	0.005
	(0.021)	(0.021)	(0.018)	(0.018)	(0.016)	(0.016)	(0.012)	(0.013)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,531	9,503	9,550	9,522	9,551	9,523	9,531	9,503
Adjusted R-squared	0.001	0.001	0.002	0.002	0.004	0.004	0.001	0.001

Panel B: Marketable Retail Trade Imbalance Trades based on # of Trades

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable	*Retail Order Imbalance (trades)*							
Window	[-3,3]		[-3,14]		[-3,30]		[-3,90]	
Operations Forward	0.024	0.024	0.062	0.063	0.057	0.057	0.047	0.048
	(0.039)	(0.039)	(0.031)	(0.036)	(0.038)	(0.038)	(0.043)	(0.043)
Operations Summary	-0.014	-0.015	-0.010	-0.001	-0.015	-0.014	0.002	0.002
	(0.018)	(0.018)	(0.016)	(0.016)	(0.012)	(0.013)	(0.010)	(0.010)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,531	9,503	9,550	9,522	9,551	9,523	9,531	9,503
Adjusted R-squared	0.001	0.001	0.002	0.002	0.008	0.008	0.015	0.015

Table 9: Visual Information, Announcement Returns, and AI Institutional Ownership

This table examines how the relation between visual information and announcement return depend on the *AI Institutional Ownership* of stocks. High and low *AI Inst. Ownership* stocks are determined by the cross-sectional annual average of *AI Inst. Ownership* in the year of presentation. The independent variables of main interest are *Operations Forward* and *Operations Summary*, which measure the forward-looking operational visual information and backward-looking operational visual information in presentation slides. We include the same set of control variables as for Table 4. All variables are defined in Appendix A. Heteroskedasticity-consistent standard errors (in parentheses) are clustered at the firm level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels (two-tailed), respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	High AI Inst. Ownership				*Low AI Inst. Ownership*			
Dependent Variables	*Car(-3, 3)*		*Car(4, 13)*		*Car(-3, 3)*		*Car(4, 13)*	
Operations Forward	0.038**	0.035**	-0.019	-0.005	0.002	0.005	-0.006	-0.02
	(0.017)	(0.017)	(0.017)	(0.018)	(0.025)	(0.025)	(0.019)	(0.020)
Operations Summary	0.006	0.008	0.005	0.005	-0.021	-0.020	0.006	-0.001
	(0.009)	(0.009)	(0.009)	(0.009)	(0.013)	(0.013)	(0.011)	(0.011)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7,773	7,752	7,771	7,750	7,780	7,763	7,775	7,750
Adjusted R-squared	0.018	0.034	0.026	0.046	0.009	0.020	0.010	0.028

Internet Appendix of "Visual Information and AI Divide: Evidence from Corporate Executive Presentations"

Table IA.1: Textual Operational Information and Market Response

This table presents results examining whether textual operational information embedded in presentation slides can predict announcement returns. The dependent variables are *Cumulative Announcement Returns* in different windows, where returns are adjusted by *Market*, *Size*, and *Book-to-Market*. The sample period is from 2006 to 2018. *Operations Forward* and *Operations Summary* measure the forward-looking operational visual information and backward-looking operational visual information in presentation slides, respectively; *Textual Operations Forward* and *Textual Operations Summary* measure corresponding textual information. We include the same set of control variables as in Table 4. Industries are defined by 2-digit SIC codes. Heteroskedasticity-consistent standard errors (in parentheses) are clustered at the firm level. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels (two-tailed), respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent Variables	*CAR(-3,3)*			*CAR(-2,2)*			*CAR(4,13)*		
Operations Forward	0.024*		0.024*	0.021*		0.023*	-0.016		-0.017
	(0.014)		(0.014)	(0.013)		(0.013)	(0.011)		(0.011)
Operations Summary	-0.009		-0.006	-0.002		-0.004	0.005		0.005
	(0.008)		(0.008)	(0.007)		(0.007)	(0.006)		(0.006)
Textual Operations Forward		0.001	-0.002		0.003	0.000		-0.015	-0.013
		(0.010)	(0.011)		(0.009)	(0.010)		(0.011)	(0.012)
Textual Operations Summary		0.002	0.002		0.004	0.003		0.002	0.002
		(0.005)	(0.005)		(0.004)	(0.004)		(0.004)	(0.004)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,326	17,143	17,273	17,326	17,143	17,273	17,275	17,141	17,271
R-squared	0.009	0.016	0.013	0.009	0.016	0.013	0.017	0.017	0.019

Table IA.2: Correlation of Visual and Textual Operational Information

This table reports the correlation of visual and textual operation information classified by deep learning model. *Operations Forward* (*Visual OF*) and *Operations Summary* (*Visual OS*) measure the forward-looking operational visual information and backward-looking operational visual information in presentation slides, respectively; *Textual Operations Forward* (*Textual OF*) and *Textual Operations Summary*(*Textual OS*) measure corresponding textual information.

	Visual OF	*Visual OS*	*Textual OF*	*Textual OS*
Visual OF	1	0.28	0.29	0.26
Visual OS	0.28	1	0.20	0.46
Textual OF	0.29	0.20	1	0.09
Textual OS	0.26	0.46	0.09	1