“Show Me!” The Informativeness of Images

Abstract

We introduce the concept of *visual informativeness* in annual reports and innovate by using machine-learning algorithms to construct visual informativeness measures. We create a novel measure of content reinforcement, representing the information content investors can extract from images, complementing and reinforcing particulars contained in textual narrative. An increase in visual prevalence and in the degree to which images convey reinforcing information is associated with greater (lower) analyst forecast accuracy (dispersion) in subsequent quarters, and lower risk. Firms increase the use of visuals when facing an exogenous drop in analyst coverage. Importantly, *visual informativeness* facilitates information assimilation.

1 Introduction

"Words! Words! Words! I'm so sick of words! ....“Show me!”

–Eliza Doolitle (MyFair Lady, Lyrics by Alan Jay Lerner and Frederick Loewe)

Information dissemination by firms reduces information processing costs (Drake, Roulstone, and Thornock, 2016; Blankespoor, 2019) and enhances price efficiency (Blankespoor, Miller, and White, 2014; Gao and Huang, 2020; Gibbons, Iliev, and Kalodimos, 2021). Improved informativeness of financial reports lessens information asymmetry and stock volatility, and boosts investment efficiency and forecasting accuracy (You and Zhang, 2009; Lawrence, 2013; Biddle, Hilary, and Verdi, 2009; Bonsall IV, Leone, Miller, and Rennekamp, 2017).

To improve information assimilation, firms have included increasingly lengthier narratives accompanied by more, graphs, charts, and images in their disclosures over time. While finance and accounting scholars have extensively researched the role of words (e.g., the *FOG* index, Li, 2008), and the frequency of graphs/charts in firms' 10-K filings, they have not examined the impact of images included in their annual reports on information assimilation and firm outcomes. This paper partially addresses this literature gap by examining whether images in particular, drive attention to and reinforce annual reports' textual information. We focus on annual reports (rather than SEC filings) because they are subject to fewer guidelines and restrictions on images and are referred to by analysts.[1]

Firms deliberately choose to disseminate annual reports widely for a reason. Presumably, they wish to convey information that is in some way distinguishable from that provided in other venues (10-Ks, 10-Qs, earning announcements, conference calls, etc). The expanded use of images, relative to other visual elements such as graphs and charts, allows for focusing our exploration on the impacts of images, an investigation that is mostly missing in the existing literature.[2] Our choice of

[1]Annual reports are far richer in graphical and image content than filings and are usually read by most stakeholders as firms post them on the investor relation section of their website. Discussion with analysts indicates that they look to annual reports to understand companies' priorities and what they want to promote. The SEC 2008 Report “Guidance on the Use of Company Web Sites” requires the format of information on firm websites to be focused on ”readability, not printability” and recognizes that “allowing companies to present data in formats different from those dictated by our forms or more technologically advanced than EDGAR may be beneficial to investors.” https://www.sec.gov/rules/interp/2008/34-58288.pdf. In fact, the SEC goes further, suggesting retail investors may take seriously: ‘boilerplate language if its presentation was visually dynamic and engaging” (SEC 2018).

[2]While other firm dissemination outlets such as conference call presentations include images, in general, they contain far fewer images. Xu (2021) finds that only 46% of conference calls include presentations, and our sampling of recent presentations indicates that of these, only 60% have image pages (24% of the conference call sample). For

annual reports as the information source for our study is also motivated by our desire to investigate images' long-term, lasting effects on financial outcomes rather than the immediate, short-term effects of visual elements, as has been emphasized in prior and ongoing research. As such, we view annual reports as a repository of textual and visual information that stakeholders occasionally consult to understand what the firms convey regarding their products and strategies.[3]

The case of American Science and Engineering, Inc., depicted in Figure 1, illustrates the richness and informativeness of image displays in annual reports. Beyond reading textual descriptions of the firm's technology in the annual report, stakeholders can glean clearer and potentially augmented and more impactful information from the images contained in the report.[4] These features can facilitate information assimilation.

We posit that images affect investors through two potentially overlapping channels. The first is attention, through which images partially relax investors' cognitive constraints (Stenning and Oberlander, 1995; Alter and Oppenheimer, 2009).[5] The second, a channel we newly explore, is *content reinforcement*, through which images reinforce concepts inherent in the narrative.[6] We provide evidence on both channels.

One challenge researchers face is systematically identifying images as distinct from other visual elements (team photos, charts/graphs, maps, and infographics). We however, are able to objectively identify distinct visual elements using machine learning algorithms and heuristic rules. Having identified these distinct visual elements, we nonetheless conjecture that annual report readers experience visuals primarily at the page level; rather than separately at the level of individual

slides with image pages, the average number of image pages is 0.8, or 2% of the presentation pages (contrasted to 7 image pages per annual report). However, while the incidence of image pages in conference calls is much lower than in annual reports, other visuals such as diagrams have been shown to impact short-term market reactions (Wu, 2021).

[3] Image pages in annual reports cannot be used for examining short-window outcomes because we cannot always ascertain the date on which the annual report was posted. Analysts, to the extent they are inclined to examine images, are more likely to consult annual reports over time due to the latter's relatively rich image content, hence giving rise to potentially longer-lasting effects. See for example, discussions in Blankespoor, deHaan, and Marinovic (2020) about the "integration costs" of processing firm disclosures.

[4] A March 3, 2020, Wall Street Journal article relays how, beyond satisfying regulatory requirements, companies engage a broad set of stakeholders by including graphics, videos, and other visual elements in their communications. https://www.wsj.com/articles/companies-find-ways-to-keep-their-annual-reports-from-being-a-bore-11583231402

[5] Scholars have acknowledged the effects of limited investor attention or processing capacity, especially when information is abundant or complex (Tversky and Kahneman, 1973; Hirshleifer, Lim, and Teoh, 2009, 2011). The psychology literature demonstrates that visuals can mitigate such effects. Experiments show that visual ease can contribute to processing fluency (Alter and Oppenheimer, 2009).

[6] As an example of lab experiments addressing the nexus between pictures and text in other contexts, see Glenberg and Langston (1992).

elements. To examine the impact of images, we therefore use *pages* that predominantly display *images*, as distinct from other visual elements, as our unit of analysis (as described in Section 3 below). We classify pages into those that do not contain visual elements (non-visual pages) and those that do (visual pages). We find that 74.3% of the reports in our sample (annual reports for S&P 1500 firms from 2002 to 2019) include visual pages; 72.5% include image pages, 40.7% have team pages, and 10.5% include pages with predominantly charts, maps, or infographics. Visual use does not seem to be concentrated in any specific industries.

We create two distinct visual measures: visual prevalence and content reinforcement. With respect to visual prevalence, while our research focuses on images, we also explore the potential impact of (and control for) other visual elements. Our main prevalence measure, *IMGC* is calculated as the number of image pages. We similarly construct the number of team pages, the union of the numbers of charts pages, maps, and infographics pages, as well as the number of pages with any visual elements. In capturing the viewer's overall exposure to visuals, these visual prevalence measures are intuitive and geared primarily towards capturing investor attention (i.e., the attention channel).

Our novel *RFC* measure is designed to capture the content reinforcement channel; it reflects how content in images enhances the as calibrated by the degree to which the algorithm's image labels are congruent to the annual report textual narrative. *RFC* is constructed following Ronen, Ronen, Zhou, and Gans (2023), by processing all image-pages through Google's Vision API, which identifies labels based on the content of the images. Specifically, *RFC* is the total number of informative image labels matching words within an annual report's narrative. We also construct the *RFC* with respect to other pertinent textual narratives produced by the firm.

We examine the determinants of firms' inclusion of image pages in annual reports. We find that greater news coverage over fiscal year t is positively associated with the number of image pages (and visual pages in general) in the subsequent annual report. Growth in total assets over the year is also positively associated with the prevalence of subsequent-year image use, seemingly to highlight expansion with visual aids. However, we detect no relationship between the firm's annual advertising expenses and image pages, suggesting firms do not merely view visuals as a marketing tool. Overall, the evidence is consistent with firms using visuals to convey information strategically.

We contribute by documenting the association between visual informativeness and analyst earnings forecast accuracy, analyst dispersion, and other firm outcomes. To our knowledge, no work to date has highlighted the impact of *images*. Notably, our work focuses on the impact of images in *annual reports*, which tend to include visuals, and which our results indicate are an additional source of information in forming analyst forecasts.[7]

We find that visual prevalence *IMGC*, the number of image pages, and content reinforcement in the annual report of fiscal year t are negatively associated with absolute forecast errors in subsequent quarters of fiscal year $t+1$, controlling for firm characteristics, disclosures, and other information dissemination. That is, analysts exhibit higher accuracy for stock A than stock B if the former is associated with more visuals or more significant content reinforcement. Using our measure of relative analyst forecast accuracy, *WAFE*, which compares forecast errors across stocks covered by each analyst in a given quarter, we detect a significant effect of visual informativeness. Notably, of the five visual element pages (image pages, team pages, maps pages, charts pages, and infographics pages), only image pages ($IMGC$) exhibit statistically and economically significant negative association with analyst forecast errors.[8] Finally, consistent with our individual analyst findings, we mainly find that *IMGC* and content reinforcement (RFC) result in lower forecast dispersion across analysts. One interpretation of this latter result is that visually conveying information causes analysts to converge in their forecasts.

Since analysts read the 10-K as well as the annual report, the nexus between the two may be relevant. We therefore construct two additional reinforcement measures, RFC_{BUS}, and RFC_{MDA} (as well as variants described below), capturing the degree to which image content reinforces the textual narrative of the firm's 10-K business description (10-K Item 1) and the firm's MD&A (10-K Item 7), respectively. Like RFC these additional measures are also associated with higher forecast accuracy. Using Natural language processing (NLP) algorithms, we also compute reinforcement

[7]10-K filings are a significant source of analysts' information set (Previts, Bricker, Robinson, and Young, 1994; Rogers and J, 1997; Gibbons et al., 2021). We contacted analysts to understand how they use annual reports in their analysis. Consistent with Gibbons et al. (2021), analysts read the 10-K financial statements as quickly as possible to update their models and publish ratings/recommendations as soon as earnings are released. But, at the same time, analysts look at the annual reports to understand companies' priorities and what they want to promote. Thus, annual reports can have long-term value in forming earnings forecasts. We explore the horizon dimension using analyst earnings forecasts over the four quarters of the subsequent fiscal year.

[8]Since our analysis is conducted at the page level, and we do not focus on the impact of individual infographics, charts, maps, or other elements, our results do not contradict those of the literature that finds that firms often use infographics and other non-image graphics that likely affect investor decisions. Nonetheless, by controlling for the effects of these visual elements in our tests, our page-level (by dominant image type on a page) results highlight the usefulness of our classification and the relevance of *images* to analyst information production.

measures based on "important sentences." The combined results support our conjecture that firms use visuals to facilitate information assimilation. We conjecture that the use of visuals can help mitigate cognitive and attention constraints analysts face, mainly when attention is limited (Hirshleifer, Levi, Lourie, and Teoh, 2019; Bourveau, Garel, Joos, and Petit-Romec, 2022) with greater benefits for more constrained analysts. Consistent with our conjecture, we find that the effect of visuals is greatest when analysts cover more stocks and multiple industries and when the text of the 10-K report is more complex.

Our combined results suggest firms use visuals to better disseminate information to investors. In our tests, we control for other information dissemination channels, such as 8-K disclosures, earnings call transcripts, management earnings guidance, and firm corporate events initiated by investor relations departments, and yet, find that visuals convey content incremental to firms' otherwise-disseminated information. However, we acknowledge that using images is an endogenous decision potentially driven by unobservables. To address endogeneity, we exploit Kelly and Ljungqvist (2012)'s brokerage closure identification strategy. We conjecture that firms are incentivized to increase visuals (images) in reports to substitute for the loss of information production once they lose coverage. Our results confirm that firms increase their use of images when facing an exogenous drop in analyst coverage. Pre- and post-event analysis supports an inference of causality.

While visuals appear to enhance the informativeness of a firm's financial report, firms may also use visuals as a marketing tool to boost their image or engender positive sentiment (hype). Our finding that visual prevalence and information reinforcement result in higher analyst accuracy lends credence to an information-enrichment story, as distinct from a sentiment-based story. Further, we find no evidence of reversals in subsequent year returns.

Indeed, our overall set of results is consistent with an information story, where visuals facilitate the assimilation of information by readers. Our tests control for firm characteristics, measures of textual readability, and other information dissemination efforts of the firm. These results do not seem to be driven by firms' stability (as captured by earnings volatility), or by analyst communication with management (captured by partitioning the sample between star and non-star analysts, or between large and small-sized brokerage firms). In addition, our identification strategy supports using visuals as a substitute for lost analyst coverage.

Our paper contributes to the established literature on information dissemination. In this study, we focus on *visual informativeness* (i.e., the use of visuals to facilitate information dissemination). as captured by content reinforcement. We show that this construct adds value above and beyond other firm information text-based measures. Visual prevalence and content–reinforcement are associated with higher analyst forecast accuracy and lower dispersion, particularly when analysts are cognitively constrained.

Of equal importance, our paper contributes to the young and growing literature that explores the use of visual information in financial settings as well as the existing literature linking visuals in other disciplines (e.g., marketing, computer science, and psychology- see Section 2 for a review of the literature). To the best of our knowledge, we are the first to explore the impact of *images* in financial reports on investors. Beyond forecast accuracy and dispersion, we also explore the impact of images on firms' risk and returns, and show that the use of images contributes to the information environment and promotes efficiency.

Finally, this paper is the first to employ two sets of novel methodologies to process visual material and tease out the distinct elements that facilitate the computation of our metrics.

The rest of the paper is organized as follows. Section 2 reviews related literature. Section 3 describes the data, explains the construction of our visual measures, and provides summary statistics. Section 4 explores the determinants of visual use in annual reports. Section 5 (6) explores the association between visual informativeness and analyst earnings forecasts (firm risk and returns). Section 7 uses brokerage closures as an identification strategy, and Section 8 concludes. Appendix A includes the variable definitions and provides additional details on the data collection process and the visual classification methodology. The Internet Appendix presents the results of using alternative visual measures, robustness and other additional tests.

2 Review of the Literature on the Use of Visual Information

We first briefly survey literature exploring the role visuals other than images play in decision-making, and then move on to discuss how researchers approached the nexus between images and investors' and others' actions. The most basic and intuitive visual aids are graphs, charts, and maps. Studies examine the visual aids' impacts on readers' financial and investment decisions in various contexts. Visual tools can increase the comprehension of information (Lusardi, Samek, Kapteyn, Glinert, Hung, and Heinberg, 2017). Household investment decisions depend on how information is displayed (Shaton, 2017). Graphical displays of Pro Forma earnings information impact even professional investors (Dilla, Janvrin, and Jeffrey, 2013). A survey experiment finds a graphic of net expected return reduces the additional (preventable) fees by up to 20% and that the visualizations' effectiveness depends on experience and familiarity with investing (Cox, de Goeij, and Van Campenhout, 2018). Christensen, Fronk, Lee, and Nelson (2023) find an *indicator* variable for infographic usage in 10-K filings to be associated in the *cross-section* with elevated post-filing stock return volatility and analyst forecast dispersion and attribute it to firms' complexity (at odds with Loughran and McDonald (2014) and Bonsall IV et al. (2017)'s finding that improved *textual* readability lowers firm risk and dispersion). In contrast, our paper which employs firm *fixed-effects* finds that visual usage is associated with lower risk and dispersion. Deng, Gao, Hu, and Zhou (2020) find an indicator of *first-time* inclusion of (any type of) graphics in a firm's annual reports to be associated with a positive return reaction and an increase in institutional holdings. Their paper differs from ours in research questions, scope, variables of interest, and methodology. Two other recent papers examine the use of charts and maps in earnings call presentations. Xu (2021) finds that firms tend to implement these more when information demand, financial statement processing costs, and operating performance are greater, and Wu (2021) finds they are used more often when information is complex and when current quarterly earnings are less persistent or fall short of analyst expectations. Wu (2021) also documents a lower post-announcement drift. Researchers have also studied the role of color in financial reports and decision-making (Chan and Park, 2015; Bazley, Cronqvist, and Mormann, 2021). When financial data are presented in red, individuals' risk preferences, expectations of future stock returns, and trading decisions are impacted (Bazley et al., 2021). Infographics accorded information a greater weight in decision-making (Bertrand and Morse, 2011).

Turning to images, we see financial reporting according them a prominent role. For example, Lee (1994), Davison and Skerratt (2007), Beattie and Jones (2000), Beattie and Jones (2008), Beattie (2014), and Davison (2014) document the use of well-known images of art masterpieces as well as commissioned artwork in firms' annual reports.[9] Lee (1994) attributes the increased use of images in financial reports to a desire to "participate in consumer engineering," wherein firms use stylized

[9] For example, British Land, Zumtobel, and WPP commissioned cartoons from Ronald Searle, Anish Kapoor, and Diego Rivera, respectively. Images of masterpieces appearing in annual reports include Vermeer's The Art of Painting (Ernst and Young's 2001 Annual Review) and Frith's Life at the Seaside (British Land Annual Report 2006).

images to induce impressions of rationality, establish the identity of the corporate personality in the minds of consumers, and influence or manipulate corporate stakeholders. UK companies with significant intangible assets were more likely to employ visual and stylistic elements in their financial reporting (Davison and Skerratt, 2007). However, Ang, Hellmann, Kanbaty, and Sood (2020) note that while graphs have been used for impression management, research on photographs (images) in financial reports is scarce in the accounting and finance literature.

A few studies do explore the relationship between the aesthetics of images and investor decisions. For example, in an experimental study, the first two pages of annual reports (more pictures, images, and more color) were found to increase the likelihood of investing in the firm (Townsend and Shu, 2010). The authors attribute this finding to increased pride of ownership in the company and a resulting increase in valuation. In different contexts, an impression of trustworthiness in photographs of potential borrowers on peer-to-peer lending sites can impact the probability of loan funding (Duarte, Siegel, and Young, 2012). Trustworthiness of clients affects auditors fees (Hsieh, Kim, Wang, and Wang, 2020) and trustworthiness and dominance of sell-side analysts are associated with lower forecast errors (Peng, Teoh, Wang, and Yan, 2022).

Pope and Sydnor (2011), Gonzalez and Loureiro (2014), and Ravina (2019) analyze how lending platforms use borrower appearance characteristics, such as race, gender, and attractiveness, in their lending decision making. Image quality can affect Airbnb booking volume (Zhang, Lee, Singh, and Srinivasan, 2017) and more videos (i.e., happy, warm, passionate) increase funding probability (Hu and Ma, 2020) Other studies have examined the effect of facial expressions, demographics, or beauty on job placement (Malik, Vir Singh, Lee, and Srinivasan, 2017); CEO compensation (Graham, Harvey, and Puri, 2017; Halford and Hsu, 2020), mutual fund performance (Ganji, Kale, and Kale, 2021); firm value (Blankespoor, Hendricks, and Miller, 2017; Halford and Hsu, 2020), and in entrepreneurial ventures (Warnick, Davis, Allison, and Anglin, 2021).

A few other contemporaneous studies examine whether and how imagery affects stock price reactions. A recent working paper by Cao, Cheng, Wang, Xia, and Yang (2022) categorizes forward-looking operational information contained in (mainly non-road show) presentation slide visuals. They find short term market reactions for firms with a high ratio of AI-equipped institutional ownership. Obaid and Pukthuanthong (2021) construct a daily market level sentiment index using news photos and find that photo pessimism predicts return reversals. Nekrasov, Teoh, and Wu

(2021) look at the existence of images in firm earnings announcement Tweets and whether the presence itself of images affects retail attention, as captured by the number of retweets and Google Search volume. Heightened attention leads to larger price reactions on earnings announcement days, but the effects reverse. Gu, Teoh, and Wu (2023) find that an investor sentiment measure they construct from StockTwits GIFs positively correlates with same-day stock returns and predicts subsequent (two-week) stock return reversals.

Obaid and Pukthuanthong (2021), Nekrasov et al. (2021), and Gu et al. (2023) focus on the effects of sentiment and attention. We complement these lines of inquiry by 1. analyzing images appearing in financial reports and 2. focusing instead on explanatory and outcome variables that are different from those used in the literature. Specifically, we consider not only the existence of images but also the type of visual content and the informativeness of images (whether reinforcing or not) as explanatory variables and analyst forecast errors, forecast dispersion, and other capital market measures as outcome variables.

As noted above, this study focuses on the ability of images to reinforce textual content in addition to their prevalence. We emphasize the objective determination and impact of images' information content rather than strictly the emotional appeal or demographic characteristics. Comparing images' information content with text-embedded content, we show how images contribute to visual informativeness and affect investors' ability to assimilate firm information.

3 Data, Visual Metrics, and Summary Statistics

In this section, we describe the annual report data we use, discuss the construction of our visual measures (Section 3.1), describe the other data sets and variables we rely on (Section 3.2), and provide summary statistics of visual measures and other firm characteristics (Section 3.3).

3.1 Annual Report Data and Visual Metrics

3.1.1 Annual Report Data

We scraped all digital annual reports available for S&P 1500 firms on Annual-Reports.com from 1989 (when data were first available) to 2019. From the 19,656 reports initially retrieved, we dropped the 1989-1992 period due to the small sample size (28 reports in total). We excluded 165 reports for which .pdf files were either broken or could not be otherwise extracted, 588 duplicate reports, 134 reports with less than five or greater than 500 pages, and 512 reports lacking the fiscal year of coverage. The resulting sample comprises 18,229 reports covering the years 1993-2019.

Table A2.1 of Appendix A.2 details this data construction process. Panel B of Table A2.1 shows the number of reports in our sample rising steadily over our sample period, potentially due to digitization and changes in the information environment. We convert each page into an 'image' file format to facilitate image processing. The 18,229 reports (before applying additional filters) comprise 2,096,775 annual report pages. We analyze data starting in 2002 (instead of 1993) because the relatively small number of firms providing digitized reports before the year 2000 raises sample selection concerns, particularly if mostly higher quality firms could apply new technologies (ahead of other firms). In addition, we require that firms have at least one news article in a given year. Since the RavenPack media coverage data starts in 2002, our final sample comprises annual reports spanning 2002 to 2019[10]

3.1.2 Images and other Visuals

Annual report pages that include visuals often combine images, graphs, charts, maps, infographics, and text. Some pages may contain many of one type of visual element, and others may provide a mix. Since the main focus of our visual investigation is on *images*, we use machine learning methods to identify them as distinct from other visual elements. Further, since our analysis is at the page level, we use similar methods to identify pages that are predominantly image pages as opposed to those that predominantly include other visual elements.

We assume readers assimilate the combined information on a page holistically and that page-level visual representation, therefore, best captures readers' focal experience; that is, we presume readers simultaneously consider and synthesize not only each element on a page, but other factors such as the size, layout, and position of the visual elements, as well as their potential interactions. Analysis at the individual image level would likely skew the importance of each image beyond the reader's experience. Indeed, since design companies offer annual report design services at the page level, not the individual level, the page-level view likely best reflects the firm's intent.[11] We,

[10]Having been provided with a lag, the 2019 data we had when we conducted the analysis is incomplete. We, therefore, exclude 2019 from our time-series statistics. Additionally, the year the annual report refers to may not correspond to the fiscal year. In such instances, we use Compustat fiscal year data. For example, Walmart's 2020 Annual Report covers the year ending January 31, 2020, corresponding to its 2019 fiscal year. Thus for this report, we use the 2019 Compustat data.

[11]Design services and templates for annual reports provide design layouts at the page level. See for exam-

therefore, conduct our investigation at the annual report *page* level.

Since report pages may contain a mix of visual elements, each varying in number, shape, positioning, and size, each visual element's impact is likely to change depending on the mix of other visuals included on the same page; a full-page-sized image can impact differently from a thumbnail appearing within a blend of different images on a page or within a mix of charts, infographics, or maps. In Figure A3.1 of Appendix A.3, the image of the cat likely captures the viewer's focal point and attention more than the small medicine dropper in the top right thumbnail image. Analysis conducted at the image level instead of the page level would have counted each thumbnail as prominently as the larger central image. It would not capture the images' relative size or, notably, their interaction and relative positioning.[12]

We combine machine learning algorithms and heuristic rules to categorize and identify visual elements and pages. We first split the 2,096,775 annual report pages in our sample into non-visual pages (those containing only text) and visual pages (containing any visual elements). For the 137,453 visual pages in the sample, we categorize the visual elements we identify on those pages into five distinct categories: images, excluding team photos (IMG); team or management photos (T); charts/graphs ($CHAR$); infographics ($INFO$); and maps ($MAPS$). Importantly, this allows us to overcome the challenge of systematically identifying images as distinct from other visual elements. Finally, using our algorithms, we categorize visual report pages by the visual element that is most prevalent on the page (image pages, team pages, charts pages, maps pages, and infographics pages). Appendix A.3 provides additional details on our visual classification methodology and the calculation of visual measures.

3.1.3 Visual Measures

We construct two distinct sets of visual measures. The first captures visual prevalence, and the second captures content reinforcement. The concept of visual prevalence is intuitive and straight-

ple Adobe instruction for design at the page level. See, for example, Adobe instruction at the page level: https://www.adobe.com/creativecloud/business/teams/resources/how-to/annual-report-design.html, and Visme, a representative software package for creating annual reports: Free Annual Report Maker - Design Reports Online — Visme.

[12] Additionally, we cannot extract individual elements from a page if the firms upload it as a combined file (.pdf or .img) of several individual image files, thus potentially underrepresenting the individual components and distorting the analysis. This also hampers the researchers' ability to identify individual elements of visual representation- for example, we cannot identify with reasonable precision or consistency pages with only images versus those with other elements within.

forward. The psychology literature suggests that visuals facilitate the flow of information and can ease cognitive constraints (e.g., Larkin and Simon, 1987; Stenning and Oberlander, 1995; Glenberg and Langston, 1992; Alter and Oppenheimer, 2009). Thus, we expect increased use of images visuals to enhance investor attention to information and facilitate information processing. Even though our emphasis in this research is on images, we explore the role other visual measures may play, either as explanatory or control variables.

Our main prevalence measure *IMGC* is calculated as the number of image pages. We also calculate other visual prevalence measures to capture (and control for) the potential impacts of non-image visual elements: *TC* (the number of team pages); *CMIC* (the union of the numbers of charts pages, maps pages, and infographics pages), and *AVC*, the union of *IMGC*, *TC*, and *CMIC*, representing the number of pages with any of these visual elements.[13]

The other measure, *RFC*, applied only to image pages, captures the content-reinforcement channel (the degree to which image information content reinforces textual narrative). We follow a three-step procedure to construct *RFC*, as in Ronen, Ronen, Zhou, and Gans (2023). First, we process each image page through the Google Vision API and analyze the algorithm-generated image labels that associate visual items with confidence levels.[14] Figure 3 presents an example of labels generated by the algorithm for an image of a woman surrounded by a pile of shoes. The top label is "footwear," with 98% confidence. Other labels pick up on the other items shown in the image, including the woman's smile, happiness, the fact that the image represents fashion, and more details regarding the specific footwear types.

Second, in order not to obfuscate the analysis with spurious word matches, we filter out image pages for which the labels are categorized as "uninformative" to obtain our final set of image pages with "informative" labels, before ascertaining whether they reinforce the textual narrative.[15] Finally, for each informative image page, we construct *RFC* as the number of informative image labels that match the annual report's text. Higher values of *RFC* represent stronger reinforcement (mapping between the image information content and the textual narrative information content).

[13] The categories of charts pages, maps pages, and infographics pages are combined to construct the *CMIC* measure because of the low incidence of their pages.

[14] https://cloud.google.com/vision.

[15] To correctly classify images, we train Google Vision on a sub-sample of images to derive a bag of words that consistently capture uninformative labels. These are used as stop labels to filter out uninformative labels based on each image page's top three labels. Appendix A.3.2 provides further detail on this process. Figure A3.2 provides examples of uninformative images.

Appendix A.3.2 details this process and lists the 100 most prevalent informative labels in our sample. These labels primarily relate to the core business operations of the company. Figure 4 presents examples of reinforcing image pages of four companies' annual reports along with their reinforcing labels, which match the text ("Vehicle" and "Motor Vehicle" for PACCAR Inc., "Furniture" for Ethan Allen Interiors, "Health Care Provider" for Teleflex Inc, and "Property" for Toll Brothers Inc).

In addition to *RFC*, we also calculate reinforcement measures with respect to the textual narrative in the business description and MD&A sections of the firm's 10-K filings: RFC_{BUS} and RFC_{MDA}, respectively, as well as three combinations of the latter: RFC_{BUSMDA}, in which we consider reinforcement of the combined text of the business and MD&A sections; $RFC_{BUSMDA+}$, in which matches are considered twice if they appear in both, and RFC_{BUSMDA_IFBOTH}; wherein matches are kept only if the image label matches a word that appears in both sections of the 10-K.

Figure 5 illustrates how image-page labels may correspond similarly or differently to these different textual narratives. For example, the first image page from the 2005 annual report of Texas Roudhouse Inc yields ten Google Vision labels. Panel B lists the labels and their probabilities, along with a breakdown of which narrative text each label matches. The word "dish" matches only the annual report text, but the word "food" matches all three narrative texts (the annual report, the business description, and the MD&A section of the firm's 10K filings).[16]

3.2 Other Data

We construct our other variables from several data sources. Stock prices, shares outstanding, and trading volume are from CRSP. Data on book value, long-term debt, total assets, sales, ROA, and advertising expenses are from Compustat. Institutional holdings are from Thomson Reuters S34 files. Data on the number of news articles for a given firm are from RavenPack, which started in 2002; we include only articles with a relevance score of 100. Finally, data on analyst coverage, analyst quarterly earnings forecasts, and analyst dispersion are from IBES.

[16] If this were the only image page in the report, the reinforcement measures for this report would be as follows: RFC= 8; RFC_{BUS}= 7; RFC_{MDA}=2.

3.3 Summary Statistics

Our sample consists of 15,477 firm-year observations from 1,363 unique firms for the period January 2002 to December 2019. To be included in the sample, a firm must be part of the S&P 1500 Index and must have been covered in the media at least once during the year. See Appendix A.2 for details regarding the data collection process.

[Table 1]

Table 1 reports the summary statistics of our visual measures and classifications. On average, the annual reports in our sample are 118 pages long. Panel A of Table 1 shows the visual prevalence and reinforcement measures distribution. The mean number of image pages (*IMGC*) per report is 5.14 (6.9% of the average number of pages), the mean number of team photos (TC) is 0.95, the mean number of pages that are dominantly charts, maps, or infographics is 0.12 and the mean *AVC* is 6.2, with a standard deviation of 8.94.

The mean number of *RFC* (text-reinforcing image labels) per report is 4.38, with a standard deviation of 6.49. Table 1 also presents the distribution of the additional reinforcement measures and their variants. Their means are 2.86, 1.81, 3.27, 3.56, and 1.14 for RFC_{BUS}, RFC_{MDA}, RFC_{BUSMDA}, $RFC_{BUSMDA+}$, and RFC_{BUSMDA_IFBOTH}, respectively.

Panel B of Table 1 reports similar statistics for the set of 11,607 annual reports that contain any visual elements (not merely text). The average number of *AVC*, *IMGC*, *TC*, and *CMIC* in the restricted sample are 8.27, 6.84, 1.26, and 0.166, with standard deviations of 9.46, 8.52, 1.90, and 0.49, respectively. The reinforcement measures are 3.81, 2.41, 4.36, 4.79, and 1.52, respectively, for *RFC*, RFC_{BUS}, RFC_{MDA}, RFC_{BUSMDA}, $RFC_{BUSMDA+}$, and RFC_{BUSMDA_IFBOTH}.

Panel C shows that the use of visual elements is pervasive across (GICS) sectors, with *AVC* ranging from 4.17 per report (Commercial Services) to 8.89 (Consumer Services), *IMGC* ranging from 3.55 (Commercial Services) to 7.57 (Consumer Services), and *TC* and *CMIC* displaying similar patterns. *RFC* ranges from 2.28 (Information Technology) to 8.04 (Consumer Staples), and, as is the case for the visual prevalence measures, the reinforcement measures are not concentrated in specific sectors.

[Table 2]

Table 2 reveals a reasonably monotonic increase in annual reports that include visual elements over time, from 297 in 2002 to a maximum of 827 in 2018. This is consistent with an overall increase in the number of reports in the sample – from 361 in 2002 to a maximum of 1,164 in 2018. On average, 74.3% of reports include visual elements. Notably, 72.5% of reports include image pages (report pages that we classify as predominantly including images). In contrast, only 40.7% of reports include team pages (those with predominantly teams) and 10.5% include CMI pages (those with predominantly charts, maps, or infographics).[17]

[Table 3]

Table 3 reports summary statistics of the main firm variables and their correlations with the various visual measures. Panel A reports the statistics of the selected firm variables. The average (median) stock market capitalization (total firm assets) is $11.66 (2.67) billion ($15.65 (3.12) billion). The average percent of institutional investors' holdings of outstanding shares is 67.5%. The percentage change in institutional holdings over the fiscal year is zero on average, with a standard deviation of 4.5%. On average, firms in our sample are covered by 129.5 news articles over the fiscal year. RavenPack's filters ensure that these articles are solely about the firm. The average ROA, cost-of-equity capital, and cost-of-debt capital are 12.3%, 11.4%, and 5.3%, respectively. On average, each firm in our sample is covered by 10 analysts.

Panel B of Table 3 reports the correlations across our visual classification measures. All measures are demeaned to capture within-firm correlations. The 0.97 correlation between *AVC* and *IMGC* confirms the importance of images as distinct from other visual elements; *AVC* and *TC* and *AVC* and *CMIC* are less correlated, at 0.55 and 0.22, respectively. Our content reinforcement measure, *RFC*, has a correlation of 0.57 with *IMGC*, suggesting that firms may use images with content reinforcement in mind. The correlations between *IMGC* and the measures capturing reinforcement of the textual narrative of the 10-K sections, RFC_{BUS}, and RFC_{MDA}, are 0.47 and 0.39, respectively. The correlation between our visual measures and *FOG* is virtually zero, suggesting that visual measures capture aspects that differ from those captured by standard text-based measures and generally improve understanding.

[17] The table reports the number of firms that include at least one report page of a specific element category in their reports, not the count of these elements in the reports.

Lastly, Panel C of Table 3 reports the correlation across the various control variables. As in Panel B, we demean the variables by firm. *LnAssets* and *LnSize* are highly correlated, and as expected, both are positively correlated with news coverage.

4 Determinants of the Use of Visual Information

In this section, we explore the determinants of visual prevalence (*IMGC*) and image content reinforcement (*RFC*). The regression specification takes the following form:

$$VIS_{j,t} = \alpha + \beta \cdot VIS_{j,t-1} + \sum_{k=1}^{K} \gamma_k \cdot X_{k,j,t} + f_j + y_t + \epsilon_{j,t}, \tag{1}$$

where the variable $VIS_{j,t}$ in the equation above is either $IMGC_{j,t}$, $RFC_{j,t}$, or $AVC_{j,t}$ derived from the annual report of firm j in year t. k denotes the specific explanatory variable; t denotes the fiscal year; and j denotes the firm. The set of explanatory variables includes the number of annual report pages (*Pages*), the natural logarithm of the total number of news articles over fiscal year t (*LnNews*), the natural logarithm of the number of firm discretionary 8-K filings (*701_801_DISCLOSURE*), the cumulative stock returns over fiscal year t (*AnnRet*), the return on assets for the fiscal year t (*ROA*), the fiscal year level of institutional holdings (*InstHold*), the annual advertising expenses normalized by annual sales (*AdvExpToSale*), the natural logarithm of the firm's assets (*LnAssets*), the natural logarithm of book-to-market ratio (*LnBM*), the daily standard deviation of returns over the fiscal year (*SdRet*), and the average daily turnover over the fiscal year (*Turnover*). We include firm fixed effects (f_j) and report–year fixed effects (y_t). To facilitate economic interpretation, we Z-Score adjust (with a mean of 0 and a standard deviation of 1) both the dependent variable and our variables of interest.[18] Note that except for *Pages*, number of annual report pages, which presumably, is decided upon simultaneously with the decision of which images to include, all other variables reflect events or decisions made throughout year t preceding the issuance of the annual report for year t. This militates against an interpretation of reverse causality.

[Table 4]

Given our focus on images as distinct from charts, maps, and other diagrammatic visuals, we use *IMGC* as our visual prevalence measure (results for *AVC*, which is the union of *IMGC*, *TC*,

[18] We exclude the stock market capitalization (*LnSize*) from these regressions because of the high correlation between *LnAssets* and *LnSize*. Replacing *LnAssets* with *LnSize* yields similar coefficients to those of *LnAssets*. However, given that we control for firm annual return, *LnAssets* better captures the changes in the firm's volume of operations.

and *CMIC*, are shown in the Internet Appendix). Panel A of Table 4 reveals a positive association between the number of news articles about the firm over the fiscal year and the number of image pages. It suggests increased media coverage reflects events included in the annual report and its associated visuals. In terms of economic significance, a one standard deviation increase in the *LnNews* results in about a 3% increase in *IMGC* standard deviation units. An increase in 8-K filings results in about a 2.2% increase in *IMGC* standard deviation units.

Annual returns and ROA are also positively correlated with the use of images; a one standard deviation increase in *AnnRet* (*ROA*) results in an increase of about 2% (3%) in the use of images. This suggests better market or accounting performance spurs the firm to enhance its use of images to highlight its success.

Visuals in annual reports may be part of the firm's advertising efforts. If so, one might expect to find a positive relationship between the firm's advertising expenses and the use of images. Our results indicate that *IMGC* is uncorrelated with adverting expenses, suggesting that image pages' prevalence does not merely reflect the firm's general marketing efforts. Also, consistent with the correlations reported in Table 3, the association between *IMGC* and *FOG* is negative but insignificant both statistically and economically.

Within columns 5-7 of Table 4, we include other firm characteristics, such as assets that reflect growth in activity. Notably, an increase in total assets is associated with an increase in *IMGC*; The association appears to be economically significant: A one standard deviation increase in assets results in an increase of 13% in *IMGC*. *LnBM* is negatively associated with *IMGC*, suggesting that higher growth (low *LnBM*) leads to higher use of visuals. *SdRet* and *Turnover* are negatively associated with *IMGC*. Results for *AVC*, our other visual prevalence measure, are reported in Table IA.1 of the Internet Appendix and are qualitatively similar to those for *IMGC*.

Panel B of Table 4 presents the determinants for the reinforcement measure, *RFC*. The results are qualitatively similar to those for *IMGC* and *AVC*. Overall, the combined results suggest that firms endeavor to convey relevant information using imagery. The following sections explore the relationship between visual use, analyst forecast errors, forecast dispersion, and a battery of firm outcomes.

5 Analysts' Earnings Forecasts and Visual Informativeness

Proceeding from a maintained hypothesis that analysts review annual reports in addition to 10-Ks – the latter typically do not contain images – we use annual reports to investigate imagery's impact on information environment variables.[19] Our focus on visual informativeness deepens our understanding of how users of annual reports assimilate information.

5.1 The Accuracy of Analysts' Quarterly Earnings Forecasts

We examine the relationship between visual informativeness and analyst quarterly earnings forecast errors by constructing a within-analyst quarterly forecast accuracy measure ($WAFE$) based on forecast errors across stocks each analyst covers in a given quarter. The measure, a variant of the *PMAFE* measure used by Clement (1999) and Jame, Johnston, Markov, and Wolfe (2016), is defined as:

$$WAFE_{i,j,q} = \frac{(AFE_{i,j,q} - \overline{AFE_{i,q}})}{\overline{AFE_{i,q}}}, \tag{2}$$

where $AFE_{i,j,q}$ is the absolute error of analyst i's forecast of firm j's earnings for fiscal quarter q of the year t+1 scaled by the stock price at the end of the previous quarter ($|Forecast - Actual|/Price_{j,q-1}$), and $\overline{AFE_{i,q}}$ is the mean absolute scaled earnings forecast error of analyst i across all stocks covered during quarter q. The regression specification takes the following form:

$$WAFE_{i,j,t+1} = \alpha + \beta \cdot VIS_{j,t} + \sum_{k=1}^{K} \gamma_k \cdot X_{k,j,t} + f_j + A_i \times y_{t+1} + \epsilon_{i,j,t+1}, \tag{3}$$

where $WAFE_{i,j,t+1}$ is the average of the four quarterly within- analyst forecast errors, ($WAFE_{i,j,q}$), for year t+1; $VIS_{j,t}$ is the selected visual measure ($IMGC_{j,t}$, $RFC_{j,t}$, or $AVC_{j,t}$), f_j is the firm fixed effect; and $A_i \times y_{t+1}$ is the analyst–year fixed effects.

We require that at least two stocks be followed by each analyst i in quarter q. We keep the analyst's most recent forecast before the earnings announcement, and require that forecasts are issued within 90 days, to avoid staleness. We control for the time lapse between the forecast date and the date of the actual earnings announcement (*DaysToEarnAnn*) – the shorter the lapse, the

[19] Conversations with an equity analyst revealed that although analysts primarily prioritize reading the 10-K financial statements as quickly as possible after earnings announcements are released to update their models and publish ratings/recommendations, they do examine the annual reports to understand companies' priorities and what they want to promote (which naturally tends to be positive). Another analyst reported that analysts do not want to be disadvantaged, since they know other analysts may examine the annual reports, they would likely follow suit.

more accurate the forecast is expected to be. We also include the textual readability *FOG* index (Lehavy, Li, and Merkley, 2011). As in Table 4, other firm control variables include: the number of annual report pages (*Pages*), the natural logarithm of the total number of news articles over fiscal year t (*LnNews*), the cumulative stock returns over fiscal year t (*AnnRet*), the return on assets for the fiscal year t (*ROA*), the fiscal year institutional holdings (*InstHold*), the annual advertising expenses normalized by annual sales (*AdvExpToSale*), the natural logarithm of the firm's assets (*LnAssets*), the natural logarithm of book-to-market ratio (*LnBM*), the daily standard deviation of returns over the fiscal year (*SdRet*), the average daily turnover over the fiscal year (*Turnover*), and the stock market capitalization as another measure of firm size (*LnSize*). We also control for analyst dispersion (*Analyst Disp*), which may affect analyst accuracy.

To ease economic interpretation, we standardize the dependent variable and the visual measures (i.e., to have a mean of zero and a standard deviation of one). As a result, the coefficients represent the effect of a one standard deviation change in X on the dependent variable in standard deviation units.

[Table 5]

Panel A of Table 5 reports results of panel regressions of analyst forecast errors on visual prevalence (measured by *IMGC*). In addition to the control variables described above, to capture the partial effect of *IMGC*, we control for the number of other visual pages (*TC* and *CMIC*).[20]

All specifications yield a negative and statistically significant coefficient, with greater visual prevalence associated with higher forecast accuracy in the subsequent year. The effect is economically significant; a one standard deviation increase in *IMGC* is associated with roughly a 2.5% increase in accuracy, measured in terms of the standard deviation of $WAFE$ (column 5). The results for *AVC* are reported in Table IA.3 of the Internet Appendix, and are qualitatively similar to those reported for *IMGC*, ranging from -2.4% to -4.2%, depending on the specification used.

Panel B reports results using the reinforcement measure (*RFC*). Again, the coefficient is negative and statistically significant across most specifications, indicating that the larger the degree to which

[20] Since most firms do not include more than one of these other visual pages per report, we use dummy indicators instead of continuous measures to capture differences between firms that use *TC* and/or *CMIC* and those that do not. Measuring *TC* and *CMIC* as continuous variables (or removing them altogether) yields qualitatively similar results.

the image content reinforces the textual narrative of the annual report, the higher the analyst forecast accuracy.

Control variables also load as expected. The coefficient of *DaysToEarnAnn* loads positively, consistent with earlier forecasts being less accurate. Institutional holdings load negatively, consistent with better governance. News and asset growth load positively, consistent with accuracy loss, potentially due to expanded operations making earnings harder to predict. Advertising expenses have a positive and significant coefficient pointing to lower forecast accuracy, suggesting perhaps that marketing expenditures such as advertising may not contribute to forecast accuracy. *FOG* has a positive and statistically significant coefficient; lower 10-K readability is associated with higher analyst forecast errors.[21] The economic significance of our visual prevalence and content reinforcement measures is comparable to that of *FOG*. Finally, *Analyst Disp* also loads positively, suggesting that stocks for which analyst forecast dispersion is higher also exhibit higher forecast errors.

5.2 Information Reinforcement with other Pertinent Narrative Text

In addition to *RFC*, we consider the extent of reinforcement of the textual narrative in the business description and MD&A sections of the firm's 10-K filing. These sections include important information about the firm's operations. Therefore, image-labels-to-text matches involving these narrative sections complement our main results by providing a gauge of reinforcement of curated and likely meaningful words.

[Table 6]

Table 6 reports the results of applying the Table 5 analysis to the RFC_{BUS}, RFC_{MDA}, and their variants' reinforcement measures. Panel A shows that each of these five measures are negatively and statistically associated with analyst forecast errors. The economic significance of the associations is at least as high as that of *RFC* (shown in Table 5), consistent with these reinforcement variants picking up words deemed important.

For robustness, Panels B-D provide results for other variations of our reinforcement measures. In Panel B, we consider measures calculated at the image page level. In Panel C, we consider measures of reinforcement of 'Important Sentences' of the 10K sections, calculated using NLP

[21] Our results on the effect of visuals on analyst accuracy are qualitatively similar when readability is alternatively measured using 10-K file size (Loughran and McDonald, 2014).

algorithms to summarize the texts of the business and MD&A sections into 10 key important sentences.[22] Panel D considers the reinforcement measures calculated at the image page level, matched to the "Important Sentences" described above. The results for all panels echo the results in Panel A: image content that reinforces textual narrative in the annual report or the 10-K filings is associated with reduced analyst forecast errors.

5.3 Visual Informativeness and Analyst Cognitive Constraints.

In this subsection, we present evidence that supports the attention/cognitive constraint channel of visual use. Like other investors, analysts have limited information processing capacity or attention (Hirshleifer et al., 2019; Bourveau et al., 2022), which we conjecture can be relaxed by visual informativeness. We consider three facets of cognitive constraints: the number of stocks analysts cover, their industry concentration, and the textual complexity of the 10-K report. Table 7 reports results for *IMGC* and *RFC*, and Table IA.4 of the Internet Appendix reports results for *AVC*. The results are qualitatively similar (for all panels) to those for *IMGC*.

[Table 7]

Panel A reports on the relationship between visuals and analyst coverage. For each subsample in each panel, the three columns correspond to columns 1, 3, and 5 of Table 5. The "High COV" ("Low COV") sub-sample comprises the highest (lowest) analyst tercile in terms of the number of stocks followed (stock coverage). As conjectured, the results for *IMGC* suggest visuals are associated with significantly smaller forecast errors for the highest tercile (those who cover more stocks, (coefficient = -0.024; t-stat= -2.36)) but not for the lowest tercile, which is roughly zero (coefficient = -0.002; t-stat= -0.11). *RFC* results are qualitatively similar, with significant results only for the highest tercile (coefficient = -0.017; t-stat = -2.18).

Panel B presents results for "High COV" analysts, ranked by their industry concentration. We conjecture that visual informativeness would have a greater impact on analysts covering a *wide* range of industries, as they may face bandwidth constraints. We calculate the maximal fraction of stocks per industry covered by each analyst. "Low Industry Concentration" ("High Industry Concentration") indicates that the analyst is in the bottom (top) tercile of industry

[22] The average number of labels matched to important sentences across the *BUSMDA* based measures is around 1, with a standard deviation of about 1.5. The correlations between *IMGC* and RFC_{BUS_IS}, RFC_{MDA_IS}, and $RFC_{BUS\_MDA\_IS}$ are 0.24, 0.16, and 0.26, respectively.

concentration. Consistent with our conjecture, results for *IMGC* indicate that visuals are associated with significantly smaller forecast errors in low industry concentrations (coefficient = -0.056; tstat = -4.53), but not in the bottom tercile (coefficient = -0.021; t- stat= -1.15). Although the coefficients for RFC are more negative for the "Low Industry Concentration' group than the "High Industry Concentration" group, the differences are minor.

Panel C explores the relationship between visual informativeness and *FOG*. The "High FOG" ("Low FOG") sub-sample is comprised of stocks that appear in the top (bottom) tercile. Consistent with our conjecture, the *IMGC* results indicate that visuals are associated with significantly smaller forecast errors for the top tercile (textually complex stocks, (coefficient = -0.055; t-stat = -2.68)), but not for the bottom tercile (coefficient = -0.014; t-stat = -1.11). Again, results for *RFC* results are consistent with those for *IMGC*.

Overall, the tests reported in Table 7 and Internet Appendix Table IA.4 support the cognitive limitation conjecture. That is, visuals and their reinforcement are most valuable when analysts face cognitive constraints in terms of the large number of stocks or industries they cover or the greater complexity of the 10-K narrative. Notably, our finding that visual informativeness is more important when *FOG* is high may be consistent with a substitution effect. When 10-K text is more obscure, analysts may resort more to the annual report (and its images) for insights.

5.4 Visual Informativeness and Firm Information Dissemination

Visual informativeness may be correlated with other information disseminated by the firm. In the above analysis, we control for firm characteristics and textual readability. In this sub-section, we further control for other firm-disseminated information events that may be correlated with visual use. We include two sets of controls.

The first set includes firm information events such as firm disclosures, investor conference events, and management earnings guidance. To capture the discretionary disclosures, we use 8-K filings (Items 7.01 and 8.01, Segal and Segal, 2016).[23] For corporate events, we rely on Bloomberg's firm

[23] While there are multiple items that can be filed with an 8-K filing, six items account for more than 96% of the cases (Ben-Rephael, Da, Easton, and Israelsen, 2022). Of those, we focus on two items that are related to firm disclosure that are somewhat subject to discretion: Item 7.01 ("Regulation FD disclosure") and Item 8.01 ("other events that are not specifically called for by Form 8-K" that the firm considers to be of importance). The other four items specifically define what triggers a filing: Item 1.01 ("entry into a material definitive agreement"), Item 2.02 ("results of operations and financial condition"), Item 5.02 ("departure/election of directors or principal officers"), and Item 5.07 ("submission of matters to a vote of security holders").

event calendar, which records all corporate activities scheduled by the firm. We use the Bloomberg function 'EVTS' and focus on 'TV/Conference/Presentation' (primarily investor conferences, including pre-scheduled press conferences), 'Analyst Marketing,' and 'Corporate Access' (consisting of firm corporate access events and analyst marketing events). We also control for the number of firm press releases throughout the year. Management earnings guidance data are obtained from I/B/E/S.

The second set of control variables includes measures extracted from year-end earnings call transcripts, which we analyze to capture any soft or additional information. We download transcripts from S&P Global and construct textual measures based on the management and Q&A transcript using the Loughran and McDonald dictionary (Loughran and McDonald, 2016), including the difference between the number of positive and negative words scaled by their sum (*SENT*), the fraction of uncertainty words (*UNC*), and strong modal words (*SMODAL*).

[Table 8]

Table 8 reports the results for *IMGC* and *RFC*. Table IA.5 of the Internet Appendix report results for *AVC*. Column 1 replicates the results of column 5 of Table 5 for reference. Columns 2 and 3 show that discretionary disclosure and corporate events are negatively and significantly associated with analyst forecast errors, consistent with both events conveying useful information. Columns 5 and 6 indicate that management earnings guidance is also negatively associated with analyst forecast errors. However, as expected, a wider guidance range has a positive, albeit non-significant coefficient. *SENT* loads negatively; one interpretation is that accuracy is lower when sentiment is negative. The results for *RFC* are qualitatively similar (Panel B), as are those for the reinforcement variants (untabulated) and *AVC*. Notably, the conclusion that visuals convey unique and valuable information is unaffected by these controls. Indeed, visuals convey content that is incremental to firms' otherwise disseminated information.

5.5 Robustness Tests

In this section, we explore possible alternative explanations for the results presented above. For parsimony, the tests described below are presented in the Internet Appendix.

Our main finding regarding analyst accuracy is that visual use is associated with higher subsequent forecast accuracy. To rule out reverse causality, we perform a "Granger Causality" test.

Specifically, we include analyst accuracy, analyst dispersion, earnings volatility, and management earnings guidance in fiscal year t as additional explanatory variables, extending the analysis conducted in Table 4. The results are reported in Table IA.10 (Internet Appendix). We find no significant relationship between any of these variables and the use of images, nor does their inclusion affect our inferences regarding the determinants discussed in Section 4.

Would "too much" visual use be detrimental? We repeat the analysis shown in Table 5, but include squared terms for our visual prevalence measures (AVC and $IMGC$). The coefficients reported in Table IA.11 are largely not significant, suggesting more numerous visual pages do not reduce accuracy.

Through easier access to management, star analysts and analysts employed at major brokerage firms may rely less on visuals. However, despite private communications amongst analysts, they may still resort to 10-Ks (Gibbons et al., 2021). It is thus likely that they would similarly resort to inspecting visuals (in the annual report) as a source of reinforcing information. We apply the analysis in Table 5 to sub-samples of analysts (star analysts vs. non-star analysts, and of brokerage firms based on the number of analysts employed). The results are reported in Table IA.12. We find significant results for both star and non-star analysts and brokerage firms' top and non-top decile sizes. That is, visual measures are associated with higher accuracy for both types of analysts and different sizes of brokerage firms.

It is possible that easier-to-forecast (stable) firms include more images in their annual reports. Also, more stable firms may lend themselves to more accurate forecasts. At the same time, more stable firms may include more images. To mitigate a concern that our results may be confounded by the coexistence of stability of the firms and the inclusion of images, we apply the analysis in Table 5 to sub-samples based on earnings volatility. The results reported in Table IA.13 suggest, interestingly, that the relationship between visuals and accuracy is stronger for firms that are harder to predict.

Finally, are analysts better at forecasting firm earnings when the outlook is positive? If so, the observed improved accuracy may not be related to the content of the images. In our last set of robustness tests, we apply the analysis in Table 5 to sub-samples based on firm ROA and firm annual returns. The results reported in Table IA.14 show the relation between visual use and accuracy to hold in both sub-samples. Notably, RFC appears to be especially important for firms that are not doing well in terms of ROA or annual returns.

5.6 Analyst Forecast Dispersion

Having established significant associations between visuals and analyst forecast accuracy, we turn our attention to analyst forecast dispersion (across analysts per covered firm). We calculate analyst dispersion as the standard deviation of the analysts' quarterly earnings forecasts normalized by the absolute mean of these forecasts and use the following model:

$$AnalystDisp_{j,t+1} = \alpha + \beta \cdot VIS_{j,t} + \sum_{k=1}^{K} \gamma_k \cdot X_{k,j,t} + f_j + y_{t+1} + \epsilon_{j,t+1}, \tag{4}$$

where $AnalystDisp_{j,t+1}$ is the average of $AnalystDisp_{j,t+1,q}$ over the four quarters for firm j in year *t+1*. *VIS* is the selected visual measure, f_j is the firm fixed effect, and y_{t+1} is the year fixed effect. We control for lagged analyst dispersion; other control variables are as in Table 5.

[Table 9]

Table 9 reports the results for *IMGC* and *RFC*. We find a negative relationship between *IMGC* and the dispersion of analyst forecasts, suggesting that visuals reduce disagreement across analysts. *AVC* results are qualitatively similar to those of *IMGC* (presented in Internet Appendix Table IA.6).

RFC is also negatively and significantly associated with analyst forecast dispersion, as are the other reinforcement measures, RFC_{BUSMDA}, $RFC_{BUSMDA+}$, and RFC_{BUSMDA_IFBOTH}, presented in Table IA.6. Our findings for the combined set of reinforcement measures support the inferences drawn from Tables 5 and 6 regarding the relevance of the content reinforcement measures to analyst output.

6 Visual Measures and Firm Risk and Return Measures

Considering that visuals are seen by investors at large, we conjecture that investors' reactions to these visuals manifest in their trading behavior, affecting risk and return. Based on the premise that stakeholders and investors in particular, refer to information disseminated by firms episodically and repeatedly over time, we focus on longer-term returns and risk measures.

We conjecture that visual measures are associated with *lower risk* in light of Loughran and McDonald (2014)'s finding of a negative association between textual readability and firm risk. To test whether visual content included in a fiscal year t annual report predicts fiscal year $t+1$'s risk measures, we focus on three dependent variables (*DEP*); the first is the standard deviation of returns (*SdRet*) over fiscal year t+1, the second is the firm's market beta (*MktBeta*) estimated using daily returns during fiscal year $t+1$; and the third is the firm's cost-of-equity (*Cost-of-Equity*) in fiscal year t+1, estimated as in Frank and Shen (2016). The regression specification takes the following form:

$$DEP_{j,t+1} = \alpha + \beta \cdot VIS_{j,t} + \sum_{k=1}^{K} \gamma_k \cdot X_{k,j,t} + f_j + y_{t+1} + \epsilon_{j,t+1}. \quad (5)$$

VIS is the visual metric of interest in the annual report of firm j in fiscal year t; k indicates the explanatory variables; t denotes the fiscal year, and j denotes the firm. The control variables are similar to those defined in Table 5 and are estimated as of the end of fiscal year t. See Appendix A.1 Table A1.1 for more details. We include firm and year fixed effects (f_j) and report firm (f_j) and year (y_t) fixed effects.

Table 10 reports the results for *IMGC*, and *RFC*. Specifications 1-2, 3-4, and 5-6 report results for *SdRet*, *MktBeta*, and *Cost-of-Equity*, respectively. Across all specifications, we find a negative association between *IMGC*, and *RFC* and the subsequent year's total risk. For example, a one standard deviation increase in *IMGC* is associated with a reduction of 1.5% in total risk (in standard deviation units). We find similar results for *MktBeta* and *Cost-of-Equity*, where a one standard deviation increase in *IMGC* (*RFC*) results in a reduction of about 1.3% (1.2%) for both *MktBeta* and *Cost-of-Equity* in standard deviation units, albeit the coefficients for *RFC* are not statistically significant. Results for *AVC* and the other reinforcement measures (RFC_{BUSMDA}, $RFC_{BUSMDA+}$, and RFC_{BUSMDA_IFBOTH}) are consistent with the results of *IMGC* and *RFC* and are presented in Internet Appendix Table IA.9.

[Table 10]

The absolute beta drop of about 0.01 (coefficient (0.018) × standard deviation (0.38)) is statistically and economically significant. The results for the cost of equity are also economically significant; the absolute drop in cost of equity is 4.5 basis points (coefficient (0.018) × standard deviation (0.025)). Overall, our findings of lower firm risk and lower cost-of-capital are consistent with firms using visuals to convey information to investors, leading to a richer informational environment.

Table IA.15 of the Appendix examines the association between year t visual measures and the firm's cumulative stock returns ($ANNRET$) in fiscal year $t+1$. we find that the associations are positive but not significant at the five percent level; notably, we observe no consequent reversals over the subsequent year. While the above effects are marginal, they point against a non-fundamental story in which firms use visuals merely as a marketing tool or where images are employed only opportunistically when performance is poor. Instead, these results, though weak, support our information-based story.

7 Identification: Brokerage Firm Closures and Visual Prevalence

Analysts play a vital information production role in the markets through their firm data synthesis and dissemination efforts. Notably, the decision to cover certain stocks is endogenous for brokerage firms.

As an identification strategy, we take advantage of Kelly and Ljungqvist (2012) 's setting and focus on terminations of sell-side analyst coverage as a result of brokerage firm closures: Between 2000-2008, 43 brokerage firms closed their research departments due to adverse changes in the economics of sell-side research, leading to 4,429 coverage terminations. Importantly, these closures were 1. well-publicized and 2. plausibly exogenous at the level of the affected stocks, as they were unrelated to individual firms' prospects. We conjecture that firms are incentivized to increase their use of visuals (images) to substitute for the loss of information production when they lose coverage.

We match the Kelly and Ljungqvist (2012) list of brokerage firms with our data and identify 23 (out of their 25) brokerage closure events that overlap (2002 to 2007) with our sample period. Control and treatment group observations during the brokerage closure year and the following year constitute a "cohort." Thus, altogether, we have six cohorts. To correct for a potential bias in staggered difference-in-difference analyses (Baker, Larcker, and Wang, 2022), we follow Gormley and Matsa (2011) and use stacked difference-in-difference regressions with cohort-firm and cohort-year fixed effects. We include first-time (but not later-treated firms) in the treatment group of each cohort. As control firms, we use non-prior-treated ones. "Pre" is the year of the brokerage firms' closures, and "post" is the post-closure year. Using propensity scores, we match on: lagged *IMGC*, *Pages*, *LnAssets*, *LnSize*, *LnBM*, *ROA*, *InstHold*, *LnNews*, *SdRet*, and *Turnover*.

We run the following model:

$$DEP_{i,j,t} = \alpha + \beta \cdot DID_{i,j,t} + \gamma_{i,j} + \omega_{i,t} + \epsilon_{i,j,t}, \tag{6}$$

where i is the cohort, j is the firm and t is the year. *DEP* is either the change in brokerage coverage in year t or the image use in the annual report of firm j in fiscal year t. *DID* is the $Treated \times Post$. All specifications include cohort-firm ($\gamma_{i,j}$) and cohort-year ($\omega_{i,t}$) fixed effects. Standard errors are clustered by firm.

[Table 11]

Panel A of Table 11 reports the results. The outcome variable in the first two columns (the first column without controls and the second column with controls) is the change in coverage ($\Delta AnalystCoverage$). We expect a drop in coverage following brokerage closures. Indeed, we find that post-brokerage-firm closure, treatment firms experienced a significant drop of 0.56 analysts on average (coefficient = -0.563; t-stat= -2.95). The second output variable in the third and fourth columns (without and with controls, respectively) is *IMGC*, which is our key variable of interest. Correspondingly, the number of image pages per annual report increased significantly by on average 1.3 pages (coefficient = 1.344; t-stat = 2.58); This increase is economically significant given the average number of image pages per annual report (5.14 for the entire sample; 6.85 for firms with any visual pages).[24]

We conduct placebo tests by running difference-in-difference regressions substituting "t+1", "t+2", "t-1", and "t-2" for the event year "t". For example, "t+1" means a shift forward of the pre- and post-period by one year. Panel B presents the results. None of the brokerage coverage drop effects on image use in the years t-1, t-2, t+1, and t+2 are statistically or economically significant, reinforcing our conjectured causal impact of reducing coverage on image use.

We also expand the sample to include up to three years after the closure year, and one additional year before; we cannot go further back because our sample begins in 2002. Including interactions of the treatment group with year indicators, Figure 6 plots the coefficient estimates and their 95% confidence intervals and confirms Panel B's finding of no significant change between the treatment

[24]We repeat the difference-in-difference analysis using analyst accuracy as an output variable. The latter is measured as the absolute earnings surprise (relative to consensus estimate) scaled by analyst dispersion (Hirshleifer, Lourie, Ruchti, and Truong, 2021). Consistent with Kelly and Ljungqvist (2012), we find a decrease in analyst accuracy in the post period (untabulated).

and the control groups in the pre-period. Our combined results support the proposition that firms increase their image use when faced with an exogenous drop in analyst coverage.

8 Conclusion

Over the past two decades, firms have increased their use of images, graphics, and other visual elements in their financial reporting. While studies have extensively explored the effect of words contained in financial reports on the firm's information environment, little is known about the determinants and effect of images on important financial outcomes.

To our knowledge, this paper is the first to explore the impact of visual informativeness – how images can improve viewer's assimilation of information in annual reports. We innovate by using machine-learning methods to tease out essential characteristics of (and categorize) the visual content we examine. Importantly, we create a novel measure of content reinforcement, representing the information content investors can extract from images, complementing and reinforcing particulars contained in the textual narrative. We show that image informativeness is an essential source of firm information dissemination.

In support of an information-based story underlying our results, we find that greater image prevalence and image content reinforcement are associated with higher (lower) analyst forecast accuracy (dispersion) and lower firm risk. Our visual measures provide consistent results across a broad set of analyzed outcomes. They are intuitive and relatively straightforward and demonstrate imagery's utility in providing information relevant to financial outcomes.

As machine learning algorithms become more advanced, we expect future research to further explore the various aspects of visual content in the information environment and financial outcomes.

Figure 1: American Science and Engineering, Inc.

This figure presents four report pages from one annual report of American Science and Engineering, Inc.

Source: Report pages from the 2008 American Science and Engineering, Inc. Annual Report.

Figure 2: Classification of Visual Elements

This Figure illustrates our five visual categories. Each report page below is identified as dominantly one of the following: images (*IMG*), team/management photos (*T*), charts (*CHAR*), maps (*MAPS*), or infographics (*INFO*).

Panel A: Images

Panel B: Teams

Panel C: Charts

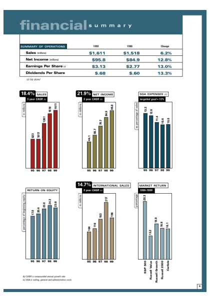

Panel D: Maps

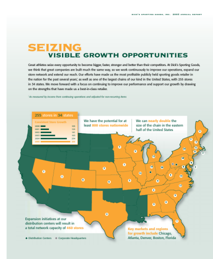

Panel E: Infographics

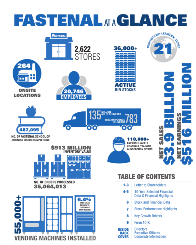

Source: (Top Left) Report page from the 2008 Coach (Now Tapestry) Annual Report; (Top Right) Report page from the 2016 Clairvest Group Inc Annual Report; (Bottom Left) Report page from the 1999 Carlisle Companies, Inc Annual Report; (Bottom Center) Report page from the 2005 DICK'S Sporting Goods Inc Annual Report; (Bottom Right) Report page from the 2015 Fastenal Co Annual Report.

Figure 3: Image-Pages and Google Vision Labels

This figure provides an example Google Vision's API Label output. The image of the woman with shoes is processed through Google Vision, and the labels are produced along with their corresponding probabilities.

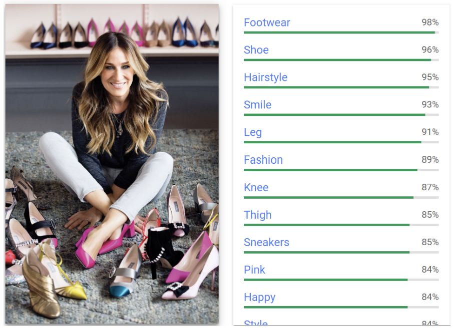

Source: https://dontmesswithtaxes.typepad.com/.a/6a00d8345157c669e20263e9633c41200b-pi.

Figure 4: Reinforcing Annual Report Image Pages

This figure provides examples of reinforcing image-pages (to the annual report textual narrative) from four annual reports in our sample. The Google Vision labels that match the text for each image-page are: For PACCAR (top left): "vehicle", "motor vehicle"; for Ethan Allen Interiors (top right): "furniture"; for Teleflex (bottom left): "health care provider; for Toll Brothers (bottom right): "property".

Source: (Top Left) Report page from the 2010 PACCAR Inc Annual Report; (Top Right) Report page from the 2007 Ethan Allen Interiors Inc Annual Report; (Bottom Left) Report page from the 2006 Teleflex Inc Annual Report; (Bottom Right) Report page from the 2013 Toll Brothers Inc Annual Report.

Figure 5: Reinforcing Image-Pages: Reinforcement to Other Textual Narrative

This figure provides examples of reinforcing image-pages. The three image-pages in Panel A are from three different annual reports in our sample. All three reinforce the textual narrative of the annual reports in which they appear, as well as both the Business (BUS) and MD&A sections of the 10-K. Panel B reports the Google Vision labels and corresponding confidence levels, along with a breakdown of which narrative text each label matches (annual report, business description, and/or the MD&A section of the firm's 10K filing).

Panel A: Reinforcing Image-pages

Source: (Top Left) Report page from the 2005 Texas Roadhouse Inc. Annual Report; (Top Center) Report page from the 2005 Mattel Inc. Annual Report; (Top Right) Report page from the 2013 Sprouts Farmers Market Inc Annual Report.

Panel B: Reinforcement Matches to Textual Narrative by Document Type

Texas Roadhouse (TXRH, NASDAQ)				
Label	Prob.	AR	BUS	MDA
dish	99%	1	0	0
food	99%	1	1	1
cuisine	99%	1	1	0
meal	95%	1	1	1
ingredient	92%	1	1	0
meat	90%	1	1	0
produce	78%	1	1	0
garnish	77%	0	0	0
carne asada	71%	0	0	0
steak	71%	1	1	0

Mattel (MAT, NASDAQ)				
Label	Prob.	AR	BUS	MDA
child	93%	1	1	1
play	83%	1	1	1
photography	78%	0	0	0
toddler	78%	0	0	0
fun	70%	1	1	0
doll	69%	1	1	1
toy	67%	1	1	1
photomontage	63%	0	0	0
collage	54%	0	0	0
child model	53%	0	0	0

Sprouts Farmers Market (SFM, NASDAQ)				
Label	Prob.	AR	BUS	MDA
ingredient	96%	1	1	0
local food	96%	1	1	0
produce	95%	1	1	1
box	94%	1	1	0
whole food	94%	0	0	0
natural foods	93%	1	1	0
marketplace	92%	1	1	0
trade	90%	1	1	1
fruit	88%	1	1	0
carton	85%	0	0	0

Figure 6: Brokerage Firm Closures - Differences by Year

This figure plots the differences between the treatment and control groups' coefficients by year. Year t is the brokerage closure event year and year $t+1$ is the first post-closure year.

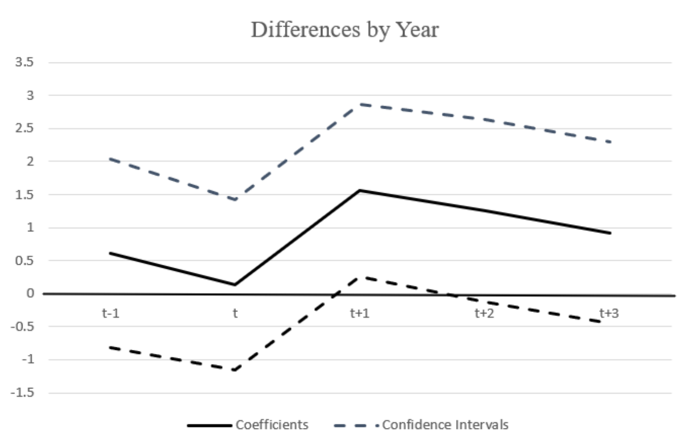

Table 1: Summary Statistics of Annual Report Pages and Visual Measures

This table reports summary statistics of annual report pages and their visual content. Panel A reports the average, standard deviation, and percentile statistics of visual prevalence measures by visual element-pages categories and of the reinforcement content measures. Panel B restricts the sample to reports that include visual elements. Panel C reports statistics for the full sample by GICS sectors. The full sample includes 15,477 firm-year observations from 1,363 unique firms and spans the years from 2002 to 2019. See Table A2.1 for the data collection process and Table C.1 for variable definitions.

Panel A: Statistics of Annual Report Pages and Visual Measures - Full Sample

	Mean	Std. Dev.	5%	10%	25%	Median	75%	90%	95%
# Report Pages	117.936	61.008	34.000	57.000	84.000	112.000	142.000	178.000	212.000
$IMGC$	5.135	7.953	0.000	0.000	0.000	3.000	7.000	12.000	16.000
TC	0.945	1.739	0.000	0.000	0.000	0.000	1.000	3.000	4.000
$CMIC$	0.124	0.428	0.000	0.000	0.000	0.000	0.000	1.000	1.000
AVC	6.204	8.940	0.000	0.000	0.000	4.000	9.000	15.000	20.000
RFC	4.378	6.486	0.000	0.000	0.000	1.000	7.000	13.000	18.000
RFC_{BUS}	2.856	4.670	0.000	0.000	0.000	0.000	4.000	9.000	13.000
RFC_{MDA}	1.805	3.286	0.000	0.000	0.000	0.000	2.000	6.000	9.000
RFC_{BUSMDA}	3.271	5.100	0.000	0.000	0.000	1.000	5.000	10.000	14.000
$RFC_{BUSMDA+}$	3.589	5.793	0.000	0.000	0.000	0.000	5.000	11.000	16.000
RFC_{BUSMDA_IFBOTH}	1.138	2.178	0.000	0.000	0.000	0.000	1.000	4.000	6.000
# of Firm-year Obs.	15,477								

Panel B: Statistics of Annual Report Pages and Visual Measures - Restricted Sample

	Mean	Std. Dev.	5%	10%	25%	Median	75%	90%	95%
# Report Pages	112.204	54.511	28.000	50.000	80.000	108.000	139.000	172.000	201.000
$IMGC$	6.847	8.522	1.000	1.000	2.000	5.000	9.000	14.000	18.000
TC	1.260	1.907	0.000	0.000	0.000	1.000	2.000	4.000	5.000
$CMIC$	0.166	0.488	0.000	0.000	0.000	0.000	0.000	1.000	1.000
AVC	8.273	9.458	1.000	1.000	3.000	6.000	11.000	17.000	22.000
RFC	5.837	6.898	0.000	0.000	0.000	4.000	9.000	15.000	20.000
RFC_{BUS}	3.809	5.045	0.000	0.000	0.000	2.000	6.000	11.000	14.000
RFC_{MDA}	2.407	3.598	0.000	0.000	0.000	1.000	4.000	7.000	10.000
RFC_{BUSMDA}	4.361	5.471	0.000	0.000	0.000	2.000	7.000	12.000	16.000
$RFC_{BUSMDA+}$	4.786	6.247	0.000	0.000	0.000	2.000	7.000	13.000	18.000
RFC_{BUSMDA_IFBOTH}	1.517	2.398	0.000	0.000	0.000	0.000	2.000	5.000	6.000
# of Firm-year Obs.	11,607								

Panel C: Statistics of Annual Report Pages and Visual Measures by GICS Sectors - Full Sample

Sector GICS code	Energy 10	Mat. 15	Ind. 20	Con. Disc. 25	Con. St. 30	Health 35	Fin. 40	Inf. Tech. 45	Com. Ser. 50	Util 55	Real Est. 60
# Report Pages(Mean)	133.54	114.65	102.41	108.92	98.90	113.00	141.21	116.61	130.10	154.85	122.16
# Report Pages(Std.Dev)	71.32	45.63	48.19	49.60	44.78	52.96	73.78	50.87	53.95	112.44	64.08
AVC(Mean)	6.25	7.61	6.62	5.93	8.89	4.91	6.90	4.40	4.17	8.16	6.20
AVC(SD)	6.53	8.65	8.07	8.25	11.04	7.66	10.55	9.45	8.60	7.12	9.01
$IMGC$(Mean)	5.28	6.40	5.53	5.13	7.57	3.91	5.18	3.75	3.55	6.69	5.24
$IMGC$(Std.Dev)	5.54	7.82	7.26	7.31	9.90	6.81	9.10	8.68	7.69	6.16	8.11
TC(Mean)	0.83	0.99	0.96	0.69	1.15	0.92	1.58	0.56	0.56	1.32	0.83
TC(Std.Dev)	1.55	1.38	1.54	1.40	1.83	1.68	2.48	1.41	1.43	1.78	1.72
$CMIC$(Mean)	0.13	0.21	0.14	0.11	0.17	0.08	0.14	0.09	0.06	0.14	0.12
$CMIC$(Std.Dev)	0.39	0.54	0.43	0.41	0.52	0.32	0.53	0.34	0.26	0.43	0.41
RFC(Mean)	4.52	5.11	5.73	5.70	8.04	2.95	3.04	2.28	2.68	5.57	3.54
RFC(Std.Dev)	6.31	6.52	7.39	7.89	9.49	4.97	4.43	3.80	4.68	6.13	5.35
RFC_{BUS}(Mean)	2.90	3.15	3.76	4.24	5.25	2.16	1.54	1.55	1.68	3.16	2.02
RFC_{BUS}(Std.Dev)	4.70	4.75	5.47	6.13	6.41	3.80	2.33	2.79	3.15	3.93	3.43
RFC_{MDA}(Mean)	1.87	1.86	2.36	2.56	3.26	1.17	1.16	0.82	1.30	2.12	1.79
RFC_{MDA}(Std.Dev)	3.21	3.41	4.00	3.90	4.80	2.39	2.09	1.77	2.59	3.13	2.99
$RFC_{BUSMDA+}$(Mean)	3.71	3.81	4.78	5.22	5.94	2.60	2.19	1.94	2.27	4.29	2.73
$RFC_{BUSMDA+}$(Std.Dev)	5.79	5.85	6.89	7.51	7.73	4.52	3.27	3.47	4.33	5.31	4.49
RFC_{BUSMDA_IFBOTH}(Mean)	1.16	1.15	1.52	1.72	1.97	0.76	0.68	0.56	0.81	1.27	0.95
RFC_{BUSMDA_IFBOTH}(Std.Dev)	2.13	2.31	2.70	2.76	3.07	1.58	1.19	1.23	1.79	1.95	1.72
# of Firm-year Obs.	656	890	2561	2346	909	1664	2148	2161	381	609	1152

Table 2: Number of Annual Reports with Visuals by Year

This table reports statistics of the number of annual reports that include visual report pages. *AR(V)* denotes reports that have pages with any visual content. *AR(I)* denotes the number of annual reports with at least one image-page, *AR(T)* denotes the number of annual reports with at least one team photos-page, *AR(CMI)* denotes the number of annual reports with at least one CMI page. The bottom row reports the time-series averages of the columns. See Table A1.1 of Appendix A.1 and Table 1 for variable and sample definitions.

FYEAR	Reports	AR(V)	%	AR(I)	%	AR(T)	%	AR(CMI)	%
2002	361	297	82.3%	289	80.1%	187	51.8%	52	14.4%
2003	534	446	83.5%	439	82.2%	275	51.5%	60	11.2%
2004	622	534	85.9%	526	84.6%	332	53.4%	81	13.0%
2005	694	584	84.1%	574	82.7%	380	54.8%	77	11.1%
2006	750	633	84.4%	615	82.0%	397	52.9%	113	15.1%
2007	804	636	79.1%	623	77.5%	385	47.9%	112	13.9%
2008	842	634	75.3%	614	72.9%	347	41.2%	90	10.7%
2009	828	619	74.8%	604	72.9%	337	40.7%	77	9.3%
2010	856	638	74.5%	623	72.8%	351	41.0%	113	13.2%
2011	864	631	73.0%	621	71.9%	301	34.8%	91	10.5%
2012	910	656	72.1%	641	70.4%	351	38.6%	103	11.3%
2013	925	651	70.4%	635	68.6%	353	38.2%	101	10.9%
2014	1003	686	68.4%	660	65.8%	379	37.8%	107	10.7%
2015	1016	697	68.6%	680	66.9%	366	36.0%	95	9.4%
2016	1129	776	68.7%	754	66.8%	381	33.7%	114	10.1%
2017	1129	776	68.7%	753	66.7%	381	33.7%	97	8.6%
2018	1164	827	71.0%	806	69.2%	366	31.4%	39	3.4%
ALL	14431	10721	74.3%	10457	72.5%	5869	40.7%	1522	10.5%

Table 3: Summary Statistics of Firm Characteristics and Correlations

This table reports summary statistics and correlations. Panel A reports the cross-sectional statistics of time series averages of the firm characteristics. Panels B and C report the correlations of our visual classification metrics, *FOG* and firm controls. All variables are demeaned by firm to capture within-firm correlations. All visual metrics are winzorised at the 99th percentile of their sample distributions. See Table A1.1 of Appendix A.1 and Table 1 for variable and sample definitions.

Panel A: Cross-Sectional Statistics

	Mean	Std. Dev.	10%	25%	Median	75%	90%
SizeInMil	11656.316	31862.225	614.562	1117.974	2671.120	8455.463	24240.240
AssetsInMil	15648.578	41756.746	460.572	1118.229	3120.500	10122.333	34012.025
BookToMarket	0.544	0.338	0.185	0.306	0.498	0.717	0.936
SdRet	0.022	0.007	0.014	0.017	0.021	0.026	0.032
Turnover	0.010	0.006	0.005	0.006	0.008	0.012	0.016
MktBeta	1.112	0.318	0.686	0.904	1.120	1.321	1.505
AnnRet	0.167	0.168	0.022	0.092	0.148	0.219	0.327
InstHold	0.675	0.161	0.454	0.579	0.700	0.792	0.853
Δ*InstHold*	-0.001	0.045	-0.041	-0.022	-0.002	0.016	0.042
#News	129.512	123.309	49.500	66.286	94.300	144.625	237.167
ROA	0.123	0.082	0.024	0.066	0.119	0.170	0.224
Cost-of-Equity	0.114	0.025	0.082	0.098	0.115	0.130	0.145
Cost-of-Debt	0.053	0.021	0.031	0.041	0.052	0.063	0.077
D/E	0.923	1.477	0.034	0.233	0.616	1.187	2.106
AdvExpToSale	0.012	0.032	0.000	0.000	0.000	0.011	0.036
AnalystsCoverage	9.985	6.743	2.909	4.796	8.042	13.964	19.817
# of firms	1,363						

Panel B: Correlations of Visual Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
(1) *AVC*	1.00										
(2) *IMGC*	0.97	1.00									
(3) *TC*	0.55	0.34	1.00								
(4) *CMIC*	0.22	0.14	0.11	1.00							
(5) *RFC*	0.58	0.57	0.30	0.10	1.00						
(6) RFC_{BUS}	0.48	0.47	0.25	0.10	0.80	1.00					
(7) RFC_{MDA}	0.40	0.39	0.21	0.08	0.71	0.72	1.00				
(8) RFC_{BUSMDA}	0.51	0.50	0.26	0.10	0.85	0.96	0.82	1.00			
(9) $RFC_{BUSMDA+}$	0.48	0.48	0.25	0.10	0.80	0.94	0.85	0.94	1.00		
(10) RFC_{BUSMDA_IFBOTH}	0.37	0.36	0.19	0.07	0.66	0.79	0.90	0.78	0.91	1.00	
(11) *FOG*	-0.02	-0.02	-0.01	-0.02	0.00	-0.00	-0.00	-0.00	-0.01	-0.00	1.00

Panel C: Correlations of Firm Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
LnNews	1.00									
AnnRet	0.01	1.00								
ROA	0.07	0.10	1.00							
AdvExpToSale	0.02	-0.03	-0.05	1.00						
LnSize	0.25	0.16	0.25	0.02	1.00					
LnBM	0.00	-0.33	-0.34	0.00	-0.47	1.00				
SdRet	0.07	0.06	-0.16	-0.01	-0.47	0.27	1.00			
Tunover	0.15	0.01	0.00	0.03	-0.09	0.05	0.42	1.00		
LnAssets	0.28	-0.10	-0.08	0.02	0.73	0.07	-0.24	-0.00	1.00	
D/E	-0.02	-0.01	-0.06	-0.00	-0.06	-0.35	0.09	0.07	0.07	1.00

Table 4: The Determinants of *IMGC* and *RFC*

This table reports results for *IMGC* and *RFC* (for completeness, Table IA.1 of the Internet Appendix reports results for *AVC*.) See Table A1.1 of Appendix A.1 and Table 1 for variable and sample definitions. All explanatory variables are measured as of end of fiscal year t. The regressions include firm and year fixed effects. Standard errors are clustered by firm and year and t-statistics are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All visual metrics are winzorised at the 99th percentile of their sample distributions. (Z) stands for a Z-Score adjustment (a mean of zero and a standard deviation of one).

Panel A: The Determinants of *IMGC*

	IMGC(Z)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
LDEP (Z)	0.432***	0.431***	0.432***	0.432***	0.429***	0.428***	0.426***
	(10.59)	(10.51)	(10.49)	(10.49)	(10.40)	(10.37)	(10.33)
Pages (Z)	-0.016	-0.015	-0.015	-0.015	-0.018	-0.016	-0.013
	(-0.88)	(-0.82)	(-0.80)	(-0.80)	(-0.98)	(-0.88)	(-0.71)
LnNews (Z)	0.037**	0.034**	0.035**	0.035**	0.019	0.016	0.030*
	(2.79)	(2.50)	(2.51)	(2.52)	(1.29)	(1.06)	(1.98)
701_801_DISCLOSURE (Z)	0.021*	0.021*	0.021*	0.021*	0.020*	0.020*	0.022*
	(1.78)	(1.82)	(1.86)	(1.85)	(1.77)	(1.77)	(1.93)
AnnRet (Z)		0.017**	0.017**	0.017**	0.020***	0.014**	0.020***
		(2.73)	(2.74)	(2.74)	(3.15)	(2.35)	(3.14)
ROA (Z)		0.041***	0.040***	0.040***	0.044***	0.034**	0.028*
		(3.06)	(2.92)	(2.91)	(3.16)	(2.45)	(2.00)
InstHold (Z)			0.002	0.002	-0.003	-0.005	-0.002
			(0.17)	(0.16)	(-0.27)	(-0.40)	(-0.18)
AdvExpToSale (Z)			-0.016	-0.016	-0.016	-0.016	-0.015
			(-0.92)	(-0.93)	(-0.87)	(-0.89)	(-0.86)
FOG(Z)				-0.004	-0.008	-0.007	-0.008
				(-0.33)	(-0.71)	(-0.70)	(-0.74)
LnAssets (Z)					0.140***	0.151***	0.132***
					(3.66)	(3.68)	(3.10)
LnBM (Z)						-0.035**	-0.030*
						(-2.27)	(-1.86)
SdRet (Z)							-0.034***
							(-2.98)
Turnover (Z)							-0.043***
							(-3.26)
Firm FE	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES
Observations	13,579	13,557	13,452	13,452	13,452	13,452	13,451
R^2	0.597	0.598	0.597	0.597	0.598	0.598	0.600

Panel B: The Determinants of *RFC*

	RFC(Z)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
LDEP	0.407***	0.407***	0.407***	0.407***	0.402***	0.402***	0.400***
	(18.00)	(18.05)	(18.09)	(18.12)	(18.03)	(17.94)	(17.74)
Pages (Z)	0.039***	0.041***	0.041***	0.041***	0.037***	0.038***	0.040***
	(3.24)	(3.34)	(3.37)	(3.36)	(2.99)	(3.10)	(3.26)
LnNews (Z)	0.035***	0.033**	0.032**	0.033**	0.011	0.009	0.017
	(2.93)	(2.71)	(2.61)	(2.66)	(0.75)	(0.62)	(1.25)
701_801_DISCLOSURE (Z)	0.012	0.013	0.012	0.012	0.011	0.011	0.012
	(1.03)	(1.05)	(1.00)	(1.00)	(0.90)	(0.92)	(1.04)
AnnRet (Z)		0.004	0.004	0.003	0.007*	0.004	0.008*
		(0.95)	(0.85)	(0.85)	(1.81)	(1.20)	(1.84)
ROA (Z)		0.024**	0.022**	0.022**	0.027**	0.022**	0.018*
		(2.49)	(2.32)	(2.32)	(2.73)	(2.20)	(1.78)
InstHold (Z)			0.007	0.006	-0.001	-0.002	-0.000
			(0.84)	(0.84)	(-0.07)	(-0.22)	(-0.02)
AdvExpToSale (Z)			-0.027	-0.027	-0.026	-0.027	-0.026
			(-1.60)	(-1.61)	(-1.54)	(-1.57)	(-1.56)
FOG(Z)				-0.003	-0.009	-0.008	-0.009
				(-0.23)	(-0.66)	(-0.64)	(-0.66)
LnAssets (Z)					0.199***	0.198***	0.187***
					(6.73)	(6.33)	(5.90)
LnBM (Z)						-0.020	-0.016
						(-1.53)	(-1.31)
SdRet (Z)							-0.020*
							(-1.89)
Turnover (Z)							-0.025***
							(-3.56)
Firm FE	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES
Firm and Year Cluster	YES	YES	YES	YES	YES	YES	YES
Observations	13,579	13,557	13,452	13,452	13,452	13,452	13,451
R^2	0.654	0.654	0.654	0.654	0.656	0.656	0.657

Table 5: Visual Prevalence, Image Content Reinforcement, and Subsequent-Year Analyst Forecast Accuracy

This table reports results from panel regressions of analyst quarterly absolute forecast errors of quarters q1–q4 in fiscal year t+1 on fiscal year t visual metrics and other explanatory variables. Panel A (B) reports results based on *IMGC* (*RFC*). Results for *AVC* are reported in Table IA.3 of the Internet Appendix. For inclusion in our analysis, we require that at least two stocks be followed by each analyst i in quarter q. We use a within-analyst quarterly forecast accuracy measure, $WAFE_{i,j,q}$. The measure is calculated as $(AFE_{i,j,q} - \overline{AFE_{j,q}})$ / $\overline{AFE_{j,q}}$, and is the absolute scaled forecast error for analyst i's forecast of firm j's earnings in quarter q of fiscal year t+1, minus the mean absolute scaled forecast error for analyst i across all the stocks the analyst follows during quarter q, divided by the mean absolute scaled forecast error of the analyst, across all stocks the analyst follows in quarter t. We average the quarterly forecast errors over fiscal year t+1. See Table A1.1 of Appendix A.1 and Table 1 for variable and sample definitions. All explanatory variables are measured as of end of fiscal year t. The regressions include firm and analyst × year fixed effects. Standard errors are clustered by analyst and year. t-statistics are reported in parentheses. Panel A also includes *TC* and *CMIC* as indicator variables. *TC* and *CMIC* coefficient estimates are reported in Table IA.2. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All visual metrics are winzorised at the 99th percentile of their sample distributions. (Z) stands for a Z-Score adjustment.

Panel A: Analyst Absolute Forecast Errors and *IMGC*

	(1)	(2)	(3)	(4)	(5)
IMGC(Z)	-0.033***	-0.034***	-0.034***	-0.024***	-0.025***
	(-3.80)	(-3.86)	(-3.85)	(-2.94)	(-3.08)
Pages	0.001***	0.001***	0.001***	0.000	0.000
	(3.32)	(3.28)	(3.26)	(1.39)	(1.36)
DaysToEarnAnn	0.001***	0.001***	0.001***	0.001***	0.001***
	(5.93)	(5.94)	(5.94)	(5.29)	(5.57)
LnNews	0.047**	0.047**	0.046**	0.080***	0.078***
	(2.16)	(2.15)	(2.11)	(2.96)	(3.01)
AnnRet	-0.185***	-0.185***	-0.185***	-0.036***	-0.032***
	(-3.37)	(-3.35)	(-3.37)	(-3.01)	(-2.68)
ROA	-2.711***	-2.694***	-2.687***	-0.317	-0.356*
	(-13.25)	(-13.09)	(-13.10)	(-1.59)	(-1.79)
insthold	-0.247***	-0.251***	-0.251***	-0.154***	-0.155***
	(-6.97)	(-7.25)	(-7.16)	(-5.03)	(-5.16)
LnAssets	-0.035	-0.036	-0.039	0.527***	0.495***
	(-1.22)	(-1.25)	(-1.37)	(16.05)	(14.18)
AdvExpToSale		1.450**	1.484**	1.509***	2.249***
		(2.33)	(2.38)	(3.35)	(2.90)
FOG(Z)			0.031***	0.022**	0.021**
			(3.55)	(2.34)	(2.29)
LnBM				-0.071***	-0.061***
				(-4.06)	(-3.67)
SdRet				6.095***	5.625***
				(3.55)	(3.44)
Turnover				-0.572	-1.064
				(-0.39)	(-0.75)
LnSize				-0.715***	-0.711***
				(-21.98)	(-21.72)
Analyst Disp					0.938***
					(4.13)
TC and *CMIC*	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES
Observations	141,867	141,831	141,831	141,829	141,759
R^2	0.490	0.490	0.490	0.520	0.524

Panel B: Analyst Absolute Forecast Errors and *RFC*

	(1)	(2)	(3)	(4)	(5)
RFC(Z)	-0.027***	-0.027***	-0.028***	-0.012	-0.012*
	(-3.57)	(-3.63)	(-3.67)	(-1.60)	(-1.72)
Pages	0.001***	0.001***	0.001***	0.000	0.000
	(3.51)	(3.48)	(3.45)	(1.43)	(1.41)
DaysToEarnAnn	0.001***	0.001***	0.001***	0.001***	0.001***
	(5.93)	(5.93)	(5.93)	(5.28)	(5.56)
LnNews	0.047**	0.047**	0.046**	0.079***	0.078***
	(2.12)	(2.11)	(2.07)	(2.96)	(3.01)
AnnRet	-0.186***	-0.186***	-0.186***	-0.037***	-0.033***
	(-3.37)	(-3.35)	(-3.37)	(-3.08)	(-2.74)
ROA	-2.721***	-2.703***	-2.696***	-0.319	-0.358*
	(-13.27)	(-13.10)	(-13.10)	(-1.60)	(-1.81)
insthold	-0.247***	-0.250***	-0.251***	-0.153***	-0.155***
	(-6.95)	(-7.23)	(-7.13)	(-5.05)	(-5.19)
LnAssets	-0.035	-0.036	-0.039	0.527***	0.495***
	(-1.20)	(-1.23)	(-1.35)	(15.85)	(14.03)
AdvExpToSale		1.425**	1.461**	1.495***	2.233***
		(2.31)	(2.36)	(3.34)	(2.90)
FOG(Z)			0.032***	0.022**	0.022**
			(3.64)	(2.35)	(2.30)
LnBM				-0.071***	-0.061***
				(-4.09)	(-3.70)
SdRet				6.169***	5.702***
				(3.58)	(3.47)
Turnover				-0.548	-1.036
				(-0.37)	(-0.73)
LnSize				-0.715***	-0.711***
				(-21.94)	(-21.65)
Analyst Disp					0.936***
					(4.10)
Firm FE	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES
Observations	141,867	141,831	141,831	141,829	141,759
R^2	0.490	0.490	0.490	0.520	0.524

Table 6: Image Content Reinforcement with Other Pertinent Textual Narrative

This table extends the results for *RFC* reported in Table 5. We replace *RFC* with additional information reinforcement measures that reinforce the narrative of the business and MD&A sections in the 10-K filings. Within each panel of this table, each row represents a distinct regression set. For parsimony, we do not report the full set of control variables below. Panel A reports results based for RFC_{BUS} and RFC_{MDA}, as well as for three variations of *RFC*; RFC_{BUSMDA}, in which we consider reinforcement to the combined text of the business and MD&A sections; $RFC_{BUSMDA+}$, in which a label that matches a word appearing in the textual narrative of both sections is counted once for each section, and RFC_{BUSMDA_IFBOTH}, which considers only label-to-text matches that appear in *both* sections. Panel B presents the results for *RFC-IP*, a variant of *RFC*, in which label-to-text matching reflects the incidence of labels per image-page (per report). We report results for $RFC\text{-}IP_{BUSMDA}$, $RFC\text{-}IP_{BUSMDA+}$ and $RFC\text{-}IP_{BUSMDA\_IFBOTH}$. In Panel C, we present results for reinforcement measures that capture label-to-text matches to (10) "Important Sentences," derived using NLP summarization tools. We report results for $RFC_{BUS\_MDA\_IS}$, $RFC_{BUSMDA+(IS)}$ and $RFC_{BUSMDA\_IFBOTH(IS)}$. Panel D presents the results for the *RFC-IP* measures based on Important Sentences ($RFC\text{-}IP_{BUSMDA(IS)}$, $RFC\text{-}IP_{BUSMDA+(IS)}$ and $RFC\text{-}IP_{BUSMDA\_IFBOTH(IS)}$). See Table A1.1 of Appendix A.1 and Table 1 for variable and sample definitions. All explanatory variables are measured as of end of fiscal year t. The regressions include firm and analyst $\times$ year fixed effects. Standard errors are clustered by analyst and year. t-statistics are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All visual metrics are winzorised at the 99th percentile of their sample distributions. (Z) stands for a Z-Score adjustment.

Panel A: *RFC* Reinforcement with Business and MD&A Sections

	(1)	(2)	(3)	(4)	(5)
$RFC_{BUS}(Z)$	-0.031***	-0.031***	-0.031***	-0.018**	-0.018**
	(-3.17)	(-3.21)	(-3.22)	(-2.17)	(-2.15)
$RFC_{MDA}(Z)$	-0.028***	-0.028***	-0.028***	-0.019**	-0.019**
	(-2.99)	(-3.02)	(-3.02)	(-2.33)	(-2.36)
$RFC_{BUSMDA}(Z)$	-0.031***	-0.031***	-0.031***	-0.018**	-0.018**
	(-3.16)	(-3.21)	(-3.21)	(-2.13)	(-2.15)
$RFC_{BUSMDA+}(Z)$	-0.030***	-0.031***	-0.031***	-0.017*	-0.017*
	(-2.96)	(-3.00)	(-3.01)	(-1.95)	(-1.97)
$RFC_{BUSMDA\_IFBOTH}(Z)$	-0.029***	-0.030***	-0.030***	-0.019**	-0.019**
	(-3.02)	(-3.04)	(-3.06)	(-2.29)	(-2.30)

Panel B: *RFC-IP* Reinforcement with Business and MD&A Sections

	(1)	(2)	(3)	(4)	(5)
$RFC\text{-}IP_{BUSMDA}(Z)$	-0.030***	-0.031***	-0.031***	-0.020**	-0.020**
	(-2.79)	(-2.83)	(-2.80)	(-2.07)	(-2.07)
$RFC\text{-}IP_{BUSMDA+}(Z)$	-0.032***	-0.033***	-0.033***	-0.022**	-0.022**
	(-2.99)	(-3.02)	(-2.99)	(-2.29)	(-2.27)
$RFC\text{-}IP_{BUSMDA\_IFBOTH}(Z)$	-0.029***	-0.030***	-0.030***	-0.021**	-0.020**
	(-3.00)	(-3.03)	(-3.01)	(-2.42)	(-2.37)

Panel C: *RFC* Reinforcement with Business and MD&A Sections - Important Sentences

	(1)	(2)	(3)	(4)	(5)
$RFC_{BUS\_MDA\_IS}(Z)$	-0.027***	-0.028***	-0.028***	-0.015***	-0.015***
	(-3.70)	(-3.75)	(-3.74)	(-2.77)	(-2.76)
$RFC_{BUSMDA+(IS)}(Z)$	-0.025***	-0.026***	-0.025***	-0.014**	-0.014**
	(-3.14)	(-3.17)	(-3.15)	(-2.34)	(-2.29)
$RFC_{BUSMDA\_IFBOTH(IS)}(Z)$	-0.010	-0.010	-0.010	-0.006	-0.005
	(-1.39)	(-1.39)	(-1.36)	(-1.09)	(-0.89)
Observations	141,867	141,831	141,831	141,829	141,759
R^2	0.490	0.490	0.490	0.520	0.524

Panel D: *RFC-IP* Reinforcement with Business and MD&A Sections - Important Sentences

	(1)	(2)	(3)	(4)	(5)
$RFC\text{-}IP_{BUSMDA(IS)}(Z)$	-0.028***	-0.028***	-0.028***	-0.017**	-0.017**
	(-3.52)	(-3.59)	(-3.59)	(-2.56)	(-2.58)
$RFC\text{-}IP_{BUSMDA+(IS)}(Z)$	-0.027***	-0.027***	-0.027***	-0.017**	-0.016**
	(-3.19)	(-3.24)	(-3.22)	(-2.41)	(-2.38)
$RFC\text{-}IP_{BUSMDA\_IFBOTH(IS)}(Z)$	-0.013*	-0.013*	-0.013	-0.009	-0.008
	(-1.68)	(-1.67)	(-1.63)	(-1.45)	(-1.27)

Table 7: Visual Measures and Analyst Cognitive Constraints

This table extends the analysis conducted in Table 5. Results are reported for *IMGC* and *RFC* based on three sub-samples: the number of stocks an analyst follows (stock coverage) (Panel A), industry concentration (Panel B), and *FOG*. (Panel C). Table IA.4 of the Internet Appendix reports results using *AVC*. For parsimony, we do not report the full set of control variables below. For each sub-sample in each panel, the three columns correspond to columns 1, 3, and 5 of Table 5. In Panel A the "High COV" ("Low COV") sub-sample is comprised of the top (bottom) analyst tercile in terms of stock coverage. Panel B presents results for analysts in the "High COV" tercile, ranked by their industry concentration. For each analyst, we calculate the max fraction of stocks in any industry they cover. "Low Industry Concentration" ("High Industry Concentration") indicates that the analyst is in the bottom (top) tercile of industry concentration. In Panel C, the "High FOG" ("Low FOG") sub-sample is comprised of stocks that appear in the top (bottom) tercile. The regressions include firm and analyst × year fixed effects. Standard errors are clustered by analyst and year. *t*-statistics are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. (Z) stands for a Z-Score adjustment. All visual metrics are winzorised at the 99th percentile of their sample distributions.

Panel A: Analyst Absolute Forecast Errors and Stock Coverage

	High COV			Low COV		
	(1)	(2)	(3)	(4)	(5)	(6)
IMGC(Z)	-0.033***	-0.033***	-0.024**	-0.011	-0.011	-0.002
	(-2.99)	(-3.01)	(-2.36)	(-0.58)	(-0.58)	(-0.11)
Firm Controls	YES	YES	YES	YES	YES	YES
TC and *CMIC*	YES	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES	YES
Observations	71,495	71,495	71,493	15,409	15,409	15,409
R^2	0.507	0.508	0.538	0.542	0.542	0.565

	High COV			Low COV		
	(1)	(2)	(3)	(4)	(5)	(6)
RFC(Z)	-0.031***	-0.032***	-0.017**	0.005	0.005	0.019
	(-3.89)	(-3.95)	(-2.18)	(0.29)	(0.30)	(1.16)
Firm Controls	YES	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES	YES
Observations	71,495	71,495	71,493	15,409	15,409	15,409
R^2	0.507	0.507	0.538	0.542	0.542	0.565

Panel B: Analyst Absolute Forecast Errors and Industry Concentration

	Low Industry Concentration			High Industry Concentration		
	(1)	(2)	(3)	(4)	(5)	(6)
IMGC(Z)	-0.066***	-0.066***	-0.056***	-0.022	-0.021	-0.021
	(-4.49)	(-4.49)	(-4.53)	(-1.11)	(-1.09)	(-1.15)
Firm Controls	YES	YES	YES	YES	YES	YES
TC and *CMIC*	YES	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES	YES
Observations	16,925	16,925	16,925	13,588	13,588	13,588
R^2	0.526	0.527	0.551	0.533	0.533	0.558

	Low Industry Concentration			High Industry Concentration		
	(1)	(2)	(3)	(4)	(5)	(6)
RFC(Z)	-0.054***	-0.054***	-0.042***	-0.042**	-0.044**	-0.035*
	(-4.18)	(-4.17)	(-3.26)	(-2.16)	(-2.26)	(-1.81)
Firm Controls	YES	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES	YES
Observations	16,925	16,925	16,925	13,588	13,588	13,588
R^2	0.526	0.526	0.550	0.533	0.533	0.558

Panel C: Analyst Absolute Forecast Errors and *FOG* - Across Stocks

	High FOG			Low FOG		
	(1)	(2)	(3)	(4)	(5)	(6)
IMGC(Z)	-0.051**	-0.051**	-0.055***	-0.013	-0.013	-0.014
	(-2.45)	(-2.47)	(-2.68)	(-0.81)	(-0.83)	(-1.11)
Firm Controls	YES	YES	YES	YES	YES	YES
TC and *CMIC*	YES	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES	YES
Observations	26,163	26,163	26,163	25,968	25,968	25,968
R^2	0.621	0.622	0.646	0.613	0.613	0.634

	High FOG			Low FOG		
	(1)	(2)	(3)	(4)	(5)	(6)
RFC(Z)	-0.070***	-0.072***	-0.040**	0.025	0.025	0.032*
	(-2.97)	(-3.00)	(-2.11)	(1.51)	(1.52)	(1.94)
Firm Controls	YES	YES	YES	YES	YES	YES
Firm FE	YES	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES	YES
Observations	26,163	26,163	26,163	25,968	25,968	25,968
R^2	0.621	0.622	0.646	0.613	0.613	0.634

Table 8: Visual Prevalence, Image Content Reinforcement, and Subsequent-Year Analyst Forecast Accuracy – Other Firm Information Dissemination Efforts

This table extends the analysis conducted in Table 5 by controlling for additional information dissemination efforts made by the firm. Panel A (B) reports results for *IMGC*(*RFC*). Results for *AVC* are reported in Table IA.5 of the Internet Appendix. In each panel, the first column reports the results from column 5 of Table 5 for reference. In all panels, *701_801_DISCLOSURE* is the log of the number of 7.01 and 8.01 items disclosed in 8K during the fiscal year. *CORPORATE EVENTS* is the log of the number of corporate events that include relevant information to investors (such as investor conferences, corporate access events, and analyst marketing events) during the fiscal year. *PRESS RELEASES* is the log of the number of firm press releases during the fiscal year. *EG* is equal to 1 if the firm issued earnings guidance during the fiscal year, and zero otherwise. *RANGE* is the management earnings guidance range normalized by the range's midpoint. Other variables are based on earnings calls transcripts. We use Loughran and McDonald's textual measures and focus on both the management and Q&A parts. The measures include the tone (*SENT*), uncertainty (*UNC*) and strong modal (*SMODAL*) of the text. The regressions include firm and analyst × year fixed effects. For parsimony, controls are not reported. Standard errors are clustered by analyst and year. *t*-statistics are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. (Z) stands for a Z-Score adjustment. All visual metrics are winzorised at the 99th percentile of their sample distributions.

Panel A: Analyst Absolute Forecast Errors and *IMGC*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IMGC(Z)	-0.025***	-0.025***	-0.025***	-0.025***	-0.024***	-0.024***	-0.024***	-0.024***
	(-3.08)	(-3.02)	(-2.99)	(-2.99)	(-2.94)	(-2.92)	(-2.96)	(-2.92)
701_801_DISCLOSURE (Z)		-0.048***	-0.048***	-0.048***	-0.048***	-0.048***	-0.047***	-0.048***
		(-3.95)	(-3.93)	(-3.93)	(-3.94)	(-3.95)	(-3.83)	(-3.93)
CORPORATE EVENTS (Z)			-0.026**	-0.025**	-0.026**	-0.026**	-0.026**	-0.024**
			(-2.51)	(-2.53)	(-2.57)	(-2.59)	(-2.52)	(-2.42)
PRESS RELEASES (Z)				-0.002	-0.002	-0.002	-0.002	-0.000
				(-0.17)	(-0.18)	(-0.16)	(-0.21)	(-0.02)
EG (Z)					-0.020*	-0.022**	-0.024**	-0.023**
					(-1.80)	(-2.10)	(-2.27)	(-2.10)
EG × RANGE (Z)						0.008	0.009	0.009
						(1.26)	(1.36)	(1.31)
MGMT SENT (Z)							-0.045***	
							(-5.34)	
MGMT UNC (Z)							0.007	
							(0.93)	
MGMT SMODAL (Z)							0.014*	
							(1.85)	
QA SENT (Z)								-0.027***
								(-3.47)
QA UNC (Z)								-0.004
								(-0.36)
QA SMODAL (Z)								0.003
								(0.32)
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
TC and *CMIC*	YES	YES	YES	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES	YES	YES	YES
Firm and Year Cluster	YES	YES	YES	YES	YES	YES	YES	YES
Observations	141,759	141,759	141,759	141,759	141,759	141,759	141,759	141,759
R^2	0.524	0.524	0.524	0.524	0.524	0.525	0.526	0.525

Panel B: Analyst Absolute Forecast Errors and *RFC*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
RFC(Z)	-0.012*	-0.012*	-0.012*	-0.012*	-0.012*	-0.012	-0.011	-0.012*
	(-1.72)	(-1.75)	(-1.73)	(-1.73)	(-1.68)	(-1.66)	(-1.56)	(-1.67)
701_801_DISCLOSURE (Z)		-0.049***	-0.049***	-0.048***	-0.049***	-0.049***	-0.048***	-0.049***
		(-3.93)	(-3.91)	(-3.91)	(-3.91)	(-3.93)	(-3.81)	(-3.90)
CORPORATE EVENTS (Z)			-0.026**	-0.026**	-0.026**	-0.026**	-0.026**	-0.025**
			(-2.51)	(-2.53)	(-2.57)	(-2.59)	(-2.53)	(-2.42)
PRESS RELEASES (Z)				-0.002	-0.002	-0.002	-0.002	-0.000
				(-0.19)	(-0.19)	(-0.17)	(-0.22)	(-0.03)
EG (Z)					-0.020*	-0.022**	-0.024**	-0.023**
					(-1.79)	(-2.09)	(-2.25)	(-2.09)
EG × RANGE (Z)						0.008	0.009	0.009
						(1.29)	(1.39)	(1.34)
MGMT SENT (Z)							-0.045***	
							(-5.32)	
MGMT UNC (Z)							0.007	
							(0.95)	
MGMT SMODAL (Z)							0.014*	
							(1.87)	
QA SENT (Z)								-0.027***
								(-3.49)
QA UNC (Z)								-0.004
								(-0.32)
QA SMODAL (Z)								0.002
								(0.24)
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
Analyst X Year FE	YES	YES	YES	YES	YES	YES	YES	YES
Firm and Year Cluster	YES	YES	YES	YES	YES	YES	YES	YES
Observations	141,759	141,759	141,759	141,759	141,759	141,759	141,759	141,759
R^2	0.524	0.524	0.524	0.524	0.524	0.524	0.525	0.525

Table 9: Visual Prevalence, Image Content Reinforcement, and Subsequent-Year Analyst Forecast Dispersion

This table reports results from panel regressions of firm dispersion of analyst earnings forecasts of quarters q1–q4 in fiscal year $t+1$ on fiscal year t visual metrics and other explanatory variables. The table reports results for *IMGC* and *RFC*. Results for *AVC* and additional reinforcement measures are reported in Table IA.6 of the Internet Appendix. $AnalystDISP_{i,j}$ is the standard deviation across the most recent analyst earnings forecasts preceding the earnings announcement date for firm i and a given quarter j, normalized the absolute value of the mean across the most recent analyst earnings forecasts. $AnalystDISP_{i,j}$ values are based on the average of the $AnalystDISP_{i,j}$ in quarters 1 to 4. As in Table 5, columns 4-6 include *TC* and *CMIC* indicators. We report their coefficients in Table IA.7. See Table A1.1 of Appendix A.1 and Table 1 for variable and sample definitions. All explanatory variables are measured as of end of fiscal year t. The regressions include firm and year fixed effects. Standard errors are clustered by firm and year, and t-statistics are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All visual metrics are winzorised at the 99th percentile of their sample distributions. (Z) stands for a Z-Score adjustment.

	(1)	(2)	(3)	(4)	(5)	(6)
IMGC(Z)	-0.024***	-0.023***	-0.015*			
	(-2.97)	(-2.95)	(-1.98)			
RFC(Z)				-0.020***	-0.020***	-0.012*
				(-2.97)	(-3.01)	(-1.85)
LagDEP	0.143***	0.141***	0.104***	0.142***	0.140***	0.103***
	(4.43)	(4.40)	(3.38)	(4.43)	(4.40)	(3.36)
Pages	0.000**	0.000**	0.000	0.000**	0.000**	0.000
	(2.38)	(2.44)	(1.55)	(2.54)	(2.60)	(1.66)
LnNews	0.059***	0.059***	0.067***	0.058***	0.058***	0.066***
	(2.80)	(2.80)	(3.18)	(2.75)	(2.75)	(3.13)
AnnRet	-0.141***	-0.143***	-0.087***	-0.141***	-0.143***	-0.086***
	(-5.22)	(-5.26)	(-3.56)	(-5.22)	(-5.26)	(-3.56)
ROA	-1.893***	-1.902***	-1.154***	-1.905***	-1.914***	-1.159***
	(-11.69)	(-11.64)	(-6.91)	(-11.90)	(-11.86)	(-6.97)
InstHold	-0.171**	-0.174**	-0.110	-0.171**	-0.174**	-0.109
	(-2.07)	(-2.12)	(-1.35)	(-2.06)	(-2.11)	(-1.34)
LnAssets	-0.024	-0.030	0.148***	-0.025	-0.031	0.148***
	(-0.83)	(-1.03)	(3.80)	(-0.86)	(-1.06)	(3.81)
AdvExpToSale		0.261	0.297		0.248	0.292
		(0.43)	(0.53)		(0.41)	(0.52)
FOG(Z)		0.022*	0.020*		0.022**	0.020*
		(1.99)	(1.91)		(2.00)	(1.90)
LnBM			-0.009			-0.009
			(-0.35)			(-0.35)
SdRet			6.932***			6.943***
			(4.39)			(4.40)
Turnover			2.886			2.961*
			(1.64)			(1.69)
LnSize			-0.213***			-0.214***
			(-4.62)			(-4.66)
701_801_DISCLOSURE (Z)			-0.039**			-0.040**
			(-2.60)			(-2.60)
CORPORATE EVENTS (Z)			0.004			0.004
			(0.25)			(0.24)
PRESS RELEASES (Z)			-0.023**			-0.023**
			(-2.03)			(-2.00)
EG (Z)			-0.018**			-0.018**
			(-2.12)			(-2.12)
EG × RANGE (Z)			0.007			0.007
			(0.63)			(0.62)
TC and *CMIC*	YES	NO	YES	NO	YES	NO
Firm FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES
Firm and Year Cluster	YES	YES	YES	YES	YES	YES
Observations	13,084	13,077	13,077	13,084	13,077	13,077
R^2	0.452	0.453	0.467	0.451	0.452	0.467

Table 10: Visual Measures and Subsequent-Year Stock Volatility, Beta and Cost-of-Equity

This table reports results from panel regressions of the firm's daily standard deviation of stock returns (*SdRet*), stock beta (*MktBeta*), and cost-of-equity capital (*Cost-of-Equity*) on fiscal year t+1 on fiscal year t visual metrics and other explanatory variables. We report results for *IMGC*, and *RFC*. As in Table 5, columns 4-6 include *TC* and *CMIC* indicators. We report their coefficients in Table IA.8. Results for *AVC* and additional reinforcement measures are reported in Table IA.9 of the Internet Appendix. See Table A1.1 of Appendix A.1 and Table 1 for variable and sample definitions. All explanatory variables are measured as of end of fiscal year t. The regressions include firm and year fixed effects. Standard errors are clustered by firm and year and t-statistics are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All visual metrics are winzorised at the 99th percentile of their sample distributions. (Z) stands for a Z-Score adjustment.

	SdRet (Z)		*MktBeta (Z)*		*Cost-of-Equity (Z)*	
	(1) *t+1*	(2) *t+1*	(3) *t+1*	(4) *t+1*	(5) *t+1*	(6) *t+1*
IMGC(Z)	-0.015**		-0.013*		-0.013*	
	(-2.26)		(-1.98)		(-2.01)	
RFC(Z)		-0.014**		-0.012		-0.011
		(-2.71)		(-1.07)		(-1.04)
LDEP	0.236***	0.236***	0.226***	0.225***	0.212***	0.212***
	(3.25)	(3.25)	(5.30)	(5.30)	(4.72)	(4.72)
Pages	0.000	0.000	-0.000	-0.000	-0.000	-0.000
	(0.68)	(0.84)	(-1.17)	(-1.07)	(-0.62)	(-0.52)
LnNews	0.021	0.021	0.070	0.069	0.083	0.082
	(1.00)	(0.98)	(1.33)	(1.32)	(1.52)	(1.52)
AnnRet	-0.039	-0.039	0.071	0.071	0.074	0.074
	(-1.01)	(-1.02)	(1.47)	(1.46)	(1.52)	(1.52)
ROA	-0.137	-0.141	-0.097	-0.101	-0.049	-0.053
	(-0.77)	(-0.79)	(-0.40)	(-0.42)	(-0.21)	(-0.23)
InstHold	-0.100**	-0.100**	0.194**	0.193**	0.194**	0.194**
	(-2.87)	(-2.89)	(2.51)	(2.51)	(2.55)	(2.54)
AdvExpToSale	-0.146	-0.159	0.913	0.898	0.899	0.885
	(-0.27)	(-0.29)	(1.52)	(1.49)	(1.45)	(1.43)
LnAssets	0.238***	0.238***	0.161**	0.161**	0.151**	0.151**
	(3.27)	(3.27)	(2.39)	(2.40)	(2.37)	(2.39)
FOG(Z)	-0.016**	-0.016**	-0.002	-0.001	-0.004	-0.004
	(-2.19)	(-2.18)	(-0.13)	(-0.11)	(-0.38)	(-0.35)
LnBM	-0.112***	-0.112***	-0.080*	-0.080*	-0.071*	-0.071*
	(-4.18)	(-4.19)	(-2.04)	(-2.04)	(-1.94)	(-1.93)
SdRet			16.544**	16.567**	16.116*	16.136*
			(2.25)	(2.25)	(2.01)	(2.01)
Turnover	3.994*	4.034*	-0.750	-0.714	-0.650	-0.612
	(1.81)	(1.83)	(-0.28)	(-0.27)	(-0.24)	(-0.23)
LnSize	-0.347***	-0.348***	-0.127**	-0.128**	-0.132**	-0.132**
	(-5.45)	(-5.46)	(-2.37)	(-2.38)	(-2.45)	(-2.46)
TDUM (Z)	-0.001		-0.005		-0.005	
	(-0.12)		(-0.64)		(-0.61)	
CMIDUM (Z)	-0.004		-0.001		-0.002	
	(-0.71)		(-0.18)		(-0.35)	
TC and *CMIC*	YES	NO	YES	NO	YES	NO
Firm FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES
Firm and Year Cluster	YES	YES	YES	YES	YES	YES
Observations	15,030	15,030	13,705	13,705	13,705	13,705
R^2	0.724	0.724	0.592	0.592	0.616	0.616

Table 11: Brokerage Mergers and Closures, Analyst Coverage, and Visual Prevalence

This table reports results from an identification strategy based on Kelly and Ljungqvist (2012)'s list of brokerage firm closures. We follow Gormley and Matsa (2011) and use stacked difference-in-difference regressions with cohort-firm and cohort-year fixed effects. For each cohort (stack) we include first-time-treated firms but not past-treated firms. As control firms we use non-prior-treated ones. "Pre" is the year of the brokerage firms' closures and "post" is post-closure year. Using propensity scores, we match on: *Pages*, *InstHold*, *LnNews*, *ROA*, *AnnRet*, *LnAssets*, *LnBM*, *SdRet*, *Turnover*, *LnSize*, and on lagged year *IMGC*. All specifications include $cohort \times firm$ and $cohort \times year$ fixed effects, and standard errors are clustered by firm. *DID* refers to $Treated \times \ Post$. In Panel A, the first two columns (without and with controls, respectively) is the change in coverage ($\Delta AnalystCoverage$) and the second output variable in the third and fourth columns (without and with controls, respectively) is *IMGC*. Panel B reports the results of a placebo test using "t+1," "t+2," "t-2," and "t-1," in columns 4, 5, 1, and 2, respectively, as if these were the event years in lieu of year 't" (the actual year of the closure). In column 3, we report the results using the actual closure year "t."

Panel A: Drop in Analyst Coverage and Increase in Image-Pages

	$\Delta AnalystCoverage$		$IMGC$	
	(1)	(2)	(3)	(4)
DID	-0.563***	-0.522***	1.344**	1.279**
	(-2.95)	(-2.75)	(2.58)	(2.47)
Cohort X Firm FE	YES	YES	YES	YES
Cohort X Year FE	YES	YES	YES	YES
Firm Controls	NO	YES	NO	YES
Observations	1,256	1,256	1,256	1,256
R^2	0.556	0.561	0.801	0.830

Panel B: Placebo Tests

	t-2	t-1	t	t+1	t+2
	(1)	(2)	(3)	(4)	(5)
DID	0.281	-0.373	1.344**	-0.525	-0.070
	(0.38)	(-0.59)	(2.58)	(-0.95)	(-0.12)
Cohort X Firm FE	YES	YES	YES	YES	YES
Cohort X Year FE	YES	YES	YES	YES	YES
Observations	504	764	1,256	1,066	812
R^2	0.829	0.786	0.801	0.802	0.828

Appendix A. Variable Definitions and Visual Classification Methodology

This appendix includes three sections:

– Appendix A.1 Includes the variable definitions;

– Appendix A.2 Describes the data collection process;

– Appendix A.3 Provides further detail on our visual classification methodology.

Appendix A.1 - Variable Definitions

Table A1.1: Variable Definitions

Variable	Definition
Visual Prevalence and Content Reinforcement Measures	
$IMGC$	For each firm, fiscal year and report, $IMGC$ is the number of image-pages (IMG) within an annual report.
TC	For each firm, fiscal year and report, TC is the number of team/management photos-pages (T) within an annual report.
$CMIC$	For each firm, fiscal year and report, $CMIC$ is the union of the numbers of charts-pages ($CHAR$), maps-pages (MAP), and infographics-pages ($INFO$) within an annual report.
AVC	The union of *IMGC*, *TC*, and *CMIC*.
RFC	For each firm, fiscal year and report, *RFC* is the number of informative labels that match words discussed in the textual narrative of the annual report.
RFC_{BUS}	For each firm, fiscal year and report, RFC_{BUS} is the number of informative labels that match words discussed in the business section of the firm 10-K report.
RFC_{MDA}	For each firm, fiscal year and report, RFC_{MDA} is the number of informative labels that match words discussed in the MD&A section of the firm 10-K report.
RFC_{BUSMDA}	For each firm, fiscal year and report, RFC_{BUSMDA} is the number of informative labels that match words discussed in the union of the business and MD&A sections of the firm 10-K report.
$RFC_{BUSMDA+}$	For each firm, fiscal year and report, $RFC_{BUSMDA+}$ is the number of informative labels that match words discussed in the business and MD&A sections of the firm 10-K report. That is, a label that appears in both sections is counted twice.
RFC_{BUSMDA_IFBOTH}	For each firm, fiscal year and report, RFC_{BUSMDA_IFBOTH} is the number of informative labels that match words discussed in *both* the business and MD&A sections of the firm 10-K report.
Textual Readability	
FOG	Gunning Fog Index (*FOG*), incorporates the number of words per sentence and the number of complex words in a document to derive a measure of the readability or syntactic complexity of firms' 10-K filings. The measure is obtained from WRDS's SEC Analytics Suite.
FILESIZE	Loughran and McDonald (2014)'s 10-K file size measure (*FILESIZE*), which is the file size (in megabytes) listed for the "complete submission text file" on EDGAR for the 10-K filing. The measure is obtained from WRDS's SEC Analytics Suite.

Variable	Definition
Firm Control Variables	
Pages	The number of pages in a given annual report.
LnNews	The natural logarithm of the total number of news articles covering the firm j in fiscal year t.
AnnRet	The 12-month cumulative stock return of firm j in fiscal year t.
ROA	The return on assets of firm j in fiscal year t.
InstHold	Aggregate institutional investor holdings based on the most recent quarter up to the end of fiscal year t. The institutional holdings data is obtained from Thomson Reuters S34 file.
Δ*InstHold*	The annual change in % institutional holdings of firm j during fiscal year t, calculated as the difference between % institutional holdings at the end of fiscal year t and the end of fiscal year t-1.
AdvExpToSale	Annual advertising expenses normalized by annual sales as in Da, Engelberg, and Gao (2011) and Lou (2014).
LnAssets	The natural logarithm of the firm's assets calculated at the end of fiscal year t.
LnSize	The natural logarithm of the firm's market capitalization calculated at the end of fiscal year t.
LnBM	The natural logarithm of the firm's book-to-market, calculated as in Fama and French (1992).
SdRet	The daily standard deviation of stock returns during fiscal year t.
Turnover	The average of the firm's daily stock turnover during fiscal year t.
D/E	The firm's debt-to-equity ratio at the end of fiscal year t.
MktBeta	Firm beta calculated using daily returns over fiscal year t.
Cost-of-Equity	The cost of equity capital (*Cost-of-Equity*) is calculated following Frank and Shen (2016). First, firm beta is calculated using daily returns over the fiscal year. Then, using the CAPM relation, the cost of equity for fiscal year t is calculated as $\textit{Cost-of-Equity} = r_f + \beta E(r_M - r_f)$. The risk-free rate, r_f, is the ten-year annualized Treasury yield from Federal Reserve economic Data (FRED). $E(r_M - r_f)$ is the historical mean of the Fama and French market excess return; that is, fiscal year t equity premium is the average of the Fama and French annualized market excess return from July 1926 to the end of fiscal year t.
AnalystsCoverage	The average number of analysts following firm j during fiscal year t.
Δ*AnalystCoverage*	The difference between the average number of analysts following firm j during fiscal year t and fiscal year *t-1*.
Earnings Volatility	The standard deviation of the firm's quarterly EPS over the previous eight quarters (fiscal years t-1 and t).

Variable	Definition
<u>Analyst Earnings Forecast Measures and Other Analyst Characteristics</u>	
AnalystDisp	The standard deviation across the most recent analyst earnings forecasts preceding the earnings announcement date for firm i and a given quarter j, normalized the absolute value of the mean across the most recent analyst earnings forecasts (obtained from IBES).
WAFE	Within-analyst quarterly forecast accuracy measure, calculated as $(AFE_{i,j,q} - \overline{AFE_{i,q}})\ /\ \overline{AFE_{i,q}}$, where $WAFE_{i,j,q}$ is the absolute forecast error for analyst i's forecast of firm j's earnings in quarter q of fiscal year t+1, minus the mean absolute forecast error for analyst i across all the stocks she follows during quarter q, divided by the mean absolute forecast error of the analyst, across all stocks she follows in quarter t. $AFE_{i,j,q}$ in turn, is the absolute error of analyst i's forecast of firm j's earnings for fiscal quarter q of the year t+1 scaled by the stock price at the end of the previous quarter $(\lvert Forecast - Actual \rvert / Price_{j,q-1})$
DaysToEarnAnn	The number of days from the forecast date to the earnings announcement date, computed for each analyst forecast in any given quarter.
Absolute Consensus Forecast Errors	The average of four absolute differences, each of which is computed as the difference between each fiscal year t's quarterly announced earnings and the corresponding quarterly forecast derived from the most recent I/B/E/S update preceding the earnings announcement, scaled by the dispersion of the analysts' forecasts (Hirshleifer et al., 2021).
Star Analyst	A dummy variable that is equal to 1 if is an analyst was named as a star analyst in fiscal year t. Information on star analysts is obtained from *Institutional Investor* Magazine's All-America Research Team rankings.
Brokerage Firm Size	The number of analysts with active earnings forecasts I/B/E/S reports as employed by the brokerage firm during fiscal year t.
<u>Other Measures</u>	
701_801_DISCLOSURE	The log of item types 7.01 and 8.01 of 8-Ks filed during the fiscal year.
CORPORATE EVENTS	We use Bloomberg's Corporate Events Calendar (the EVTS function) to obtain information about scheduled corporate events such as investor conference events, corporate access events, and analyst marketing events. *CORPORATE EVENTS* is the log of the number of corporate events during the fiscal year. The data are available from 2010.
PRESS RELEASES	The log of the number of firm press releases during the fiscal year. The data are obtained from RavenPack's press-release file.
EG	A dummy variable that is equal to 1 if the firm issued earnings guidance during the fiscal year, and zero otherwise.
RANGE	The management earnings guidance range normalized by the range's midpoint.

Appendix A2: Data Collection Process

This Appendix describes the data collection process (Panel A) and provides time-series statistics (Panel B) for our annual report data.

Table A2.1: Data Collection Process

This table describes the annual report data construction process. We downloaded and analyzed all digitally available reports for S&P 1500 firms trading in the United States (with a matched PERMNO) between 1989 and 2019. We applied filters to ensure data integrity and availability in arriving at the final sample reported in Table 1 as outlined below (Panel A). Panel B reports the time series statistics of firms' annual reports containing visual elements (AV) from 1993 to 2019. *# REPORTS* is the number of firms with annual reports. *# AV REPORTS* is the number of annual reports with visual elements. *# PAGES* is the total number of annual report pages across all reports in a given year. See Table A1.1 of Appendix A.1 for variable definitions. Any Visual (AV) pages are those for which any visual elements can be detected on the report page, where visual elements have an image size of at least 100K or vividness of at least 100.

Panel A: Data Filtering Process

Procedure Description	Sample
Firm annual reports collected for S&P 1500 firms between 1989 and 2019	19,656
Less reports from 1989 to 1992	28
Less reports that broken and cannot be opened	165
Less reports that are duplicated	588
Less reports with >= 500 or <= 5 pages	134
Less reports with no fiscal year identified	512
Final sample 1993-2019 before additional filters	18,229
Keeping the sample between 2002 and 2019	16,861
Keeping firms with media coverage	15,477

Panel B: Time-Series Statistics Before Additional Restrictions

FYEAR	*# REPORTS*	# AV	%	*# Report Pages*
1993	21	7	33.3%	2,651
1994	32	14	43.8%	3,142
1995	44	19	43.2%	4,545
1996	65	31	47.7%	7,351
1997	104	59	56.7%	8,571
1998	157	100	63.7%	10,944
1999	252	188	74.6%	15,514
2000	338	272	80.5%	21,913
2001	402	325	80.8%	26,106
2002	482	402	83.4%	37,173
2003	578	485	83.9%	46,377
2004	663	569	85.8%	58,879
2005	741	624	84.2%	67,822
2006	802	675	84.2%	78,498
2007	857	682	79.6%	90,241
2008	902	681	75.5%	102,974
2009	889	671	75.5%	101,743
2010	924	687	74.4%	110,377
2011	942	694	73.7%	113,142
2012	997	715	71.7%	124,834
2013	1,025	722	70.4%	131,987
2014	1,072	731	68.2%	139,259
2015	1,122	772	68.8%	146,080
2016	1,221	837	68.6%	161,373
2017	1,218	840	69.0%	160,807
2018	1,250	891	71.3%	169,465
2019	1,129	760	67.3%	155,007
ALL	18,229	12,438	68.2%	2,096,775

Appendix A3: Visual Classification Methodology

In this Appendix, we describe how we classify annual report pages based on their visual content into the categories depicted in Figure 2 (sub-section A3.1) and the additional steps to construct the *RFC* measures (sub-section A3.2).

A3.1. Classification of Visual Pages using Machine Learning Tools

To identify which report pages contain visual elements and which do not, we first construct a training sample of report pages with visual elements (AV "any visual") and those without. Based on this sample, we trained a TensorFlow classification mode (based on Google Brain open-source machine learning and AI software library for training and inference of deep neural networks) to do binary classification on all report pages. We combined human judgment with color composition and page size to construct a representative training sample. Specifically, for each report page, we extracted the 16 basic HTML colors (black, white, grey, red, yellow, etc.) and calculated the color composition/distribution. If the primary colors (over 90% of pixels) are black, white, and grey, the page is not classified as a visual page; if more colors are contained in the page, the page is classified as visual. We first convert each report page into image format to calculate page size and then calculate its physical file size. Visual pages contain colors, different shapes, styles, etc., and are thus more likely to be larger in file size. We combined the objective information obtained from color and file size with subjective judgment processes to get our initial validation sample, which yielded a 96% accuracy rate.

For those pages classified as *AV* we combined artificial intelligence and a rule-based system to classify visual annual report pages in our sample pages into our 5 five distinct predefined hierarchical categories: image pages, team pages; charts pages; maps-pages; and infographics pages. Instead of manually selecting training samples for each visual element category, we rely on the following simple process: First, we process the pages through the Google Vision API and identify which labels typify each of the five visual categories. We then classify each page into the corresponding categories using these labels. For example, if a page yields the label "map", this page is assigned to the MAPS group.

This initial classification process yields a pool of page candidates for each category. Based on these candidate pools, we construct a representative training sample to train a classification model. If a page contains visual content of more than one category, we categorize it by the dominant visual category that best describes its visual content. Figure 2 provides an example of visual pages that have been classified into the five categories. When there are mixed visuals on a page, the classification is based on the dominant category. Finally, we train TensorFlow to classify as above.

Image pages (*IMG*) were categorized with an accuracy rate of about 97%. The remaining four visual categories were classified with an accuracy rate of approximately 71%. To increase the accuracy rate, we augmented our algorithms using Google Vision and heuristic rules to increase the accuracy of the other categories. For example, if one of the top three Google Vision labels for a visual page contains the word "map" or "maps", then this page is classified as a maps page. This combined approach improved classification accuracy rates of map-pages, charts-pages, and teams-pages to approximately 86% and of infographics-pages to roughly 78%.

Infographics-pages typically contain a broad mix of text, fonts, colors, numbers, icons, small graphs, shapes, and/or photos. They are, therefore difficult to identify using machine learning methods and are often subject to misclassification error. We therefore rely on Google Tesseract Optical Character Recognition to capture the location, size, and style of textual elements. Combining the information from these last two steps, we then use a rule set to reclassify those misclassified infographics, increasing infographic classification accuracy from 78% to 85%, comparable to the accuracy rate of the other visual categories.

After validating these procedures, we applied them to the remaining visual pages in our sample. Finally, we removed pages that could not be classified in one of the five visual categories with a certain threshold (50% by default). Most of these were textual pages printed on a non-white page (for example, a blue background with text written in black).

A3.2. Construction of the *RFC* Measure

To construct *RFC*, our measure of reinforcement of image content to narrative text, we follow a three-step procedure, as in Ronen et al. (2023). We process each of the image pages through Google Vision and analyze the algorithm-generated image labels that associate visual items with confidence levels. We first filter out images that are uninformative as follows: 1. we derive stop labels by training Google Vision on a sub-sample of images to identify a bag of words that consistently capture uninformative labels, which are: “adaptation,” “aqua,” “atmosphere,” “atmospheric phenomenon,” “azure,” “background,” “beige,” “black,” “black-and-white,” “blue,” “brown,” “circle,” “cobalt blue,” “colorfulness,” “daytime,” “design,” “diagram,” “document,” “drawing,” “ecoregion,” “electric blue,” “floor plan,” “font,” “fractal art,” “graphic,” “graphic design,” “graphics,” “gray,” “green,” “grey,” “illustration,” “leaf,” “light,” “line,” “line art,” “liquid,” “logo,” “magenta,” “map,” “maroon,” “material property,” “music,” “orange,” “organism,” “paper,” “paper product,” “parallel,” “pattern,” “pie chart,” “pink,” “plan,” “plot,” “poster,” “purple,” “rectangle,” “red,” “schematic,” “screenshot,” “sky,” “slope,” “space,” “square,” “teal,” “technical drawing,” “text,” “triangle,” “turquoise,” “violet,” “water,” “watermark,” “wave,” “white,” “world,” and “yellow”; 2. suppose any of the top three labels for an image corresponds to a stop label. In that case, the image is filtered out as “uninformative.” Figure A3.2 presents examples of representative uninformative image pages from four different annual reports. The top three labels produced by Google Vision are listed in descending order of confidence. In each of these examples, all three of the leading labels are stop-labels.

Finally, we process all the labels for informative image pages and calculate the number of informative image labels per image page that match the text of the annual report. For each report, we then construct (RFC) by summing the matches from all image pages in the report. RFC is calculated without double counting of label-to-text matches, while $RFC\text{-}IP$ adds label-to-text matches across all image pages in a report. Other variants considered in the paper examine the reinforcement of labels to alternative narrative text sources, such as the business description and MD&A sections of 10-K filings, for RFC_{BUS} and RFC_{MDA}, respectively.

Table A3.1 provides a list of the 100 most prevalent informative labels in our sample, and Figure 4 presents examples of reinforcing image pages and their reinforcing labels.

In contrast, Figure A3.3 provides examples of informative image pages that are not reinforcing (no label-to-text matches exist). In the top right image, for example, the labels generated for the 2004 Hanmi Financial Corporation Annual Report are “bowed string instrument”, “cellist”, “cello”, “classical music”, “musical instrument”, “musician”, “recital”, “string instrument”, and “violin family.” None of these match any of the words in the textual narrative of the 2004 Annual Report. and the image page is thereby classified as non-reinforcing.

Table A3.1: Reinforcing Labels – Examples

This table lists the 100 most prevalent informative labels (top three for each image page), ranked by the number of years they appear in our sample. "AnnualReport" reports the top 100 labels that match the annual report, "BUS" reports the top 100 labels that match the business description section of the firm's 10-K filing, and "MD&A" reports the top 100 labels that match the MD&A section for the firm's 10-K filing.

Rank	AnnualReport	BUS	MD&A	Rank	AnnualReport	BUS	MD&A
1	aircraft	aircraft	building	51	cuisine	dish	book
2	engineering	engineering	advertising	52	dish	floor	cap
3	advertising	advertising	brand	53	pipe	pipe	community
4	architecture	architecture	car	54	skin	skin	locomotive
5	brand	brand	city	55	website	website	metal
6	brochure	brochure	engineering	56	airline	airline	produce
7	building	building	event	57	art	denim	research
8	car	car	food	58	boat	gas	clothing
9	city	city	infrastructure	59	denim	shoe	company
10	clothing	clothing	nature	60	gas	soil	field
11	event	event	number	61	jeans	art	road
12	food	food	property	62	sharing	beauty	boat
13	hand	hand	service	63	shoe	boat	bottle
14	human	human	table	64	soil	bottle	construction equipment
15	infrastructure	infrastructure	transport	65	arm	company	electricity
16	metal	metal	vehicle	66	beauty	dress	history
17	motor vehicle	motor vehicle	aircraft	67	book	footwear	meal
18	nature	nature	furniture	68	bottle	jeans	pipe
19	number	number	hand	69	cap	style	recipe
20	plant	plant	plant	70	company	construction equipment	sharing
21	property	property	fashion	71	dress	electricity	soil
22	publication	publication	human	72	electricity	farm	steel
23	recipe	room	industry	73	field	field	denim
24	room	service	joint	74	footwear	fruit	home
25	service	tire	machine	75	head	head	shelf
26	table	transport	retail	76	history	home	shoe
27	tire	vehicle	room	77	home	interior design	winter
28	transport	wheel	website	78	metropolitan area	meal	agriculture
29	vehicle	wood	airline	79	style	metropolitan area	alcohol
30	wheel	customer	architecture	80	air travel	road	asphalt
31	wood	drink	aviation	81	construction equipment	sharing	bridge
32	aviation	electronics	construction	82	farm	trade	flooring
33	customer	face	customer	83	fruit	arm	hospital
34	drink	fashion	face	84	interior design	asphalt	metropolitan area
35	electronics	furniture	people	85	meal	cabinetry	motor vehicle
36	face	industry	airplane	86	research	cap	motorcycle
37	fashion	ingredient	brochure	87	road	locomotive	mountain
38	floor	joint	footwear	88	shelf	research	railway
39	furniture	machine	ingredient	89	trade	crane	restaurant
40	industry	people	publication	90	asphalt	flooring	trade
41	ingredient	produce	tire	91	cabinetry	heat	travel
42	joint	recipe	air travel	92	crane	history	beer
43	machine	retail	child	93	lip	hospital	collection
44	people	table	dish	94	locomotive	lawn	cuisine
45	produce	airplane	drink	95	media	media	dog
46	retail	aviation	electronics	96	railway	motorcycle	drilling rig
47	airplane	child	farm	97	travel	shelf	lawn
48	child	community	floor	98	winter	steel	ship
49	community	construction	gas	99	agriculture	truck	tractor
50	construction	cuisine	wood	100	bridge	chair	watch