Seeing is Believing: Annual Report Enhanced Visuals and Stock Returns

Wesley (Xiaohu) Deng, Lei Gao, Bo Hu, Guofu Zhou*

First Draft: January, 2020
Current Version: March, 2025

Abstract

Why do companies employ visually enhanced annual reports that appear redundant to the plain 10-K filings? We develop a model to explain this phenomenon. Using a comprehensive dataset and advanced AI-powered analysis of visual content, we report the first evidence that firms earn positive abnormal returns after newly adding visuals to their annual reports, accompanied by a surge in institutional investor attention and holdings. Notably, while firms using info-graphics (e.g., tables, charts, plots) do not exhibit abnormal returns, the effect is more pronounced among those employing other illustrative and photographic images—particularly visuals emphasizing innovation and technological advancements. Furthermore, firms adopting R&D-focused visuals experience approximately 17% rise in patents granted, followed by a 16% increase in new product launches—a pattern consistent with communicating forward-looking innovation capabilities ahead of tangible outcomes. These findings align with our theory that companies strategically utilize visuals to mitigate investor inattention and better convey nuanced information to fundamental investors.

Keywords: *Strategic Visual Disclosure, Investor Inattention, Fundamental Investors, Technology Advancement and Innovation, Big Data, AI*

* Deng is from University of New South Wales UNSW Business School. Gao and Hu are from George Mason University Costello College of Business. Zhou is from Washington University in St. Louis Olin Business School. Correspondence: Guofu Zhou, email: zhou@wustl.edu; phone: (314) 935-6384. We appreciate helpful comments from Daniel Collins, Tony Cookson, John Dooley, Richard Evans, Gerard Hoberg, Edith Hotchkiss, Yael Hochberg, Benjamin Kessler, Pete Kyle, Charles Lee, Kai Li, Christian Lundblad, Andrey Malenko, Ron Masulis, Maximilian Muhn, Gordon Phillips, Kuntara Pukthuanthong, Tarun Ramadorai, Lukas Schmid, Noah Stoffman, Yongxian Tan, Vladimir Volkov, Hang Wang, Julie Wu, Zhuo Zhong, as well as seminar participants at Cheung Kong Graduate School of Business, CUKH-Shenzhen, Curtin University, Dongbei University of Finance and Economics, George Mason University, Jinan University, Macquarie University, Syracuse University, University of Bath, University of Liverpool, Washington University in St. Louis, Xi'an Jiaotong-Liverpool University, and participants at 2021 International FinTech Conference, FIRN Asset Pricing Meeting, and Pacific Basin Finance, Economics, Accounting and Management Conference, 2022 CICF, 2022 SFS Cavalcade Asia-Pacific, 2023 Bretton Woods Accounting and Finance Ski Conference, 2024 Australasian Finance & Banking Conference, and 2024 European Finance Association Annual Meeting.

Companies find ways to keep their annual reports from being a bore. Videos, graphics and other elements can provide information in ways that engage and go beyond what regulators request. *Wall Street Journal*, 3/3/2020

As perhaps the most important financial document for investors, the annual report of a public firm has two versions: one is the mandatory electronic 10-K filing with the U.S. Securities and Exchange Commission (SEC), the other is the user-friendly annual report available on the firm website or delivered to investors. Under the SEC Regulation S-T, there should be no material differences between the two versions, although the latter version can be customized by including large, colorful graphics which are generally absent in the plain 10-Ks.[1] In fact, more than half of the firms listed in the S&P 1500 are releasing such *visual annual reports*. This raises an economic puzzle: why are so many companies producing aesthetically designed annual reports if they are informationally redundant to the plain 10-Ks?

At first glance, the visuals in annual reports may appear cosmetic or marketing-driven.[2] If that were the case, they would have little impact on stock valuations, aside from potentially triggering short-term sentiment and noise trading. However, it is also plausible that these visuals convey information not readily apparent in plain 10-K filings.[3] Beyond capturing investor attention, visuals can enhance comprehension of subtle but valuable business information. This paper presents the first theoretical and empirical study on the informational value embedded in the visuals of firms' annual reports. Research on the role of visuals in financial disclosures remains in its early stages. Here, we examine the visuals in firm annual reports (arguably the most fundamental corporate disclosures) and provide a rational framework to explain their emergence and significance.

Theoretically, we develop a simple model that explains this phenomenon and generates three testable predictions. Our model builds on the seminal framework of Kyle (1985), but relaxes a key assumption that all public information is fully incorporated by all market participants. Instead, we assume that certain aspects of a firm's value or payoff are too subtle to be effectively conveyed and, as a result, escape the attention of most traders, consistent with the literature on investor inattention (Hirshleifer and Teoh (2003), Peng and Xiong (2006), Barber and Odean (2008), Da, Engelberg, and Gao (2011), Lou (2014), and Alldredge and Cicero (2015)). Consequently, the stock may become undervalued, motivating

[1] See Section I for the institutional background and Internet Appendix A for more details.

[2] Regulators may view them as "glossy"; see, for example, https://www.sec.gov/files/reada10k.pdf.

[3] Regulation S-T Rule 304(b) requires firms to disclose in their electronic filings any visual elements contained in their annual reports delivered to investors. However, "a good faith effort" is acceptable, which may leave room for extra subtle information embedded in graphic annual reports but not visible in 10-Ks.

the firm manager to put more efforts for more effective communication. Attentive investors may notice this shift and investigate the subtle information, referred to as alpha. However, it remains uncertain who will discover and capitalize on it. Based on rational expectations, the manager acts only when the expected alpha exceeds an endogenous threshold.[4] On average, this leads to positive abnormal returns following the manager's efforts, as reflected in the increased use of visuals in annual reports. Conversely, when the expected alpha falls below the threshold, the manager opts for inaction, drawing no attention and prompting no research or discovery of alpha. Our model predicts that if speculative trading does not immediately react to the manager's information-driven actions and if fundamental investors gradually identify and process the extra information, stock price reactions will be delayed. Thus, the presence of a delayed anomaly would indicate that the return predictability associated with increased visual content in annual reports is novel to the market.

Empirically, we collect annual reports of 5,008 companies from 1994 to 2022, and develop an OpenAI-powered algorithm (along with Salesforce's model) to quantify and classify the visual elements within each report. After processing over 8 million image files, our program identifies around 2 million images with strong visual impacts based on color and size, and further categorizes 1.15 million images using their descriptive captions. This allows us to construct various metrics of visual usage in annual reports. After merging with accounting and return data and removing firm-year observations from the regulated industries, our final sample includes 277,980 categorized image files, corresponding to 9,704 firm-year observations for 1,390 firms from 1994 to 2022. Consistent with our model predictions, we find that firms experience positive abnormal returns following the release of *visually enhanced annual reports*, accompanied by increased attention and holdings of institutional investors. Investors' reaction is gradual, with a delay of approximately one month, suggesting that investors require time to recognize and process the extra information associated with visually enhanced reports.[5]

We further test two mechanisms through which visuals may influence investors and hence stock prices: reinforcement of quantitative information and attention-enhanced comprehension of qualitative information. We employ an advanced AI technique to classify images into two categories: (1) info-graphics, which include tables, plots, charts, and diagrams that visually reinforce numerical or textual information; and (2) illustrative and photographic visuals

[4]This is also related to the literature on voluntary disclosures, following the seminal works of Grossman (1981), Milgrom (1981), Verrecchia (1983, 1990, 2001), Diamond (1985), and Diamond and Verrecchia (1991). Our model focuses on the relationship between firm communication efforts and abnormal stock returns.

[5]The delay can have various sources. For example, the information may be *soft* so that it requires extra time for analysis and verification (Liberti and Petersen (2019)). Hierarchical organizations may also delay investment decisions, especially for fundamental investors (Evans, Rohleder, Tentesch, and Wilkens (2022)).

(IP visuals), such as photos of people and scenes, and exhibitions of products, technologies, and services.[6] This classification enables us to track annual changes in firms' reporting format as to whether they issue new IP visuals and info-graphics. Our results indicate that firms issuing new visual report IP visuals is associated with positive abnormal returns, whereas a rise in info-graphics has no significant effect. Additionally, firms that expand their use of IP visuals attract significantly greater institutional attention, while firms adopting new info-graphics do not.

We then analyze which aspects of IP visuals drive the effect and find that it is largely driven by visuals emphasizing innovation, technological advancements, biotech developments, and other R&D-related domains. These areas are inherently complex and demand significant effort from investors to fully understand. Our empirical results align with our predictions and with the regulatory guidelines that encourage the use of visuals in corporate disclosures. Notably, firms incorporating new R&D-related visuals tend to file and receive more patents in subsequent years, followed by an increase in new product introductions. This finding suggests that those type of visuals provide signals of future changes in firm fundamentals. Such signals are often subtle and suggestive, requiring human interpretation rather than pure machine analysis. Thus, the observed abnormal returns may reflect the value of fundamental analysis, despite the growing dominance of quantitative trading in modern markets.

In summary, this paper demonstrates the predictive power of a non-textual, highly unstructured dataset — the visuals embedded in firms' annual reports. These visual elements turn out to be a valuable source of financial information overlooked by most market participants. Consistent with our model, our empirical results suggest that R&D-related visuals have impact on institutional investors who may receive valuable hints for their fundamental analysis and investment decisions. The effect is positive and persistent, supported by real activities positively affecting firm fundamentals.

Cognitive studies have shown that visual representations are more engaging and accessible to humans than textual formats.[7] Visuals offer unique advantages: they capture attention quickly, enhance memory, improve understanding, and transcend language barriers. These benefits explain the growing popularity of data visualization in various contexts and fields.[8] As noted by Hawkins and Hawkins (1986), "*the average investor is confused and overburdened by the volume and detail of the annual report, unable to comprehend its complexities and*

[6] Logos are also part of IP visuals, but since they typically do not convey meaningful information about firms' fundamentals, we exclude them from our main analysis.

[7] See Larkin and Simon (1987), Kosslyn (1989), Pinker (1990), Stenning and Oberlander (1995), and Roth and Bowen (2003), among others.

[8] See, for example, Tufte (1985), Friel, Curcio, and Bright (2001), Kastellec and Leoni (2007), Healey and Enns (2011), Smiciklas (2012), Schwabish (2014), among others.

jargon, and unable to use all of the information contained in the annual report." Visuals solve this problem by quickly engaging the audience and organizing their attention to the key takeaways. In the business world, it is often difficult to quantify or textualize qualitative information. The advantages of a new technology or product may not be fully captured by technical descriptions or numerical data but may be more effectively conveyed through demonstrations of its novel applications. For example, Nvidia's graphics processing units (GPUs) enabled a major breakthrough in HIV research. Perilla and Klaus (2017) simulated the HIV capsid on a supercomputer with 3880 GPU-accelerated nodes and published their finding in the *Nature* magazine. This news was weaved into Nvidia's annual report, together with *Nature*'s cover page that features a 3D art of the HIV particle. This visual can quickly direct investor attention to an impressive application. Similar visuals can convey implicit messages that exert subtle influences, allowing managers to hint at hidden strengths or intangible assets not well reflected in textual statements.

For reasons mentioned above, we postulate that when firm managers intend to convey subtle information beyond words and numbers, they will preferentially use visual presentations. We keep taking Nvidia as an illustrative example. This company began releasing image-rich annual reports since 2010. Figure 1 presents Nvidia's compound abnormal returns[9] alongside the numbers of visuals in its annual reports from 2006 to 2020. Obviously, the surge in visuals in Nvidia's reports preceded its accelerated equity growth in recent years.

It may be surprising that stock prices do not respond immediately to the emergence of IP visuals. This indicates that those visuals do not directly convey extra *material* information, consistent with Regulation S-T Rule 304 that any material information presented in graphics must be fully disclosed in the 10-K filings as well. If visuals conveyed new material information, we would expect prompt, market-wide responses which would destroy the predictive power of visuals in annual reports. Nonetheless, visuals may deliver extra *non-material* information[10] that hints about new fundamentals. Fundamental investors do not bet on hints. They do research and aggregate various sources of data (e.g., firm investments) which may enable them to develop nonpublic material information and adjust their investments.

According to our analysis, the extra signal from visually enhanced annual reports tend to be subtle, suggestive, and uncertain. These characteristics complicate both managerial communication and investor interpretation. Investors may lose interest in firms with inef-

[9]This is estimated by the CAPM model. We use the CRSP data of Nvidia's betas (with the Nasdaq index as the benchmark portfolio) to calculate its annual excess returns which are then compounded over time.

[10]By definition, non-material information does not directly affect security prices but may provide some hints or insights to investors regarding the possible future performance of the security. See the SEC's Rule of *Selective Disclosure and Insider Trading* at https://www.sec.gov/rules/final/33-7881.htm and a related discussion at https://www.sec.gov/rules/proposed/s73199/zeikel1.htm.

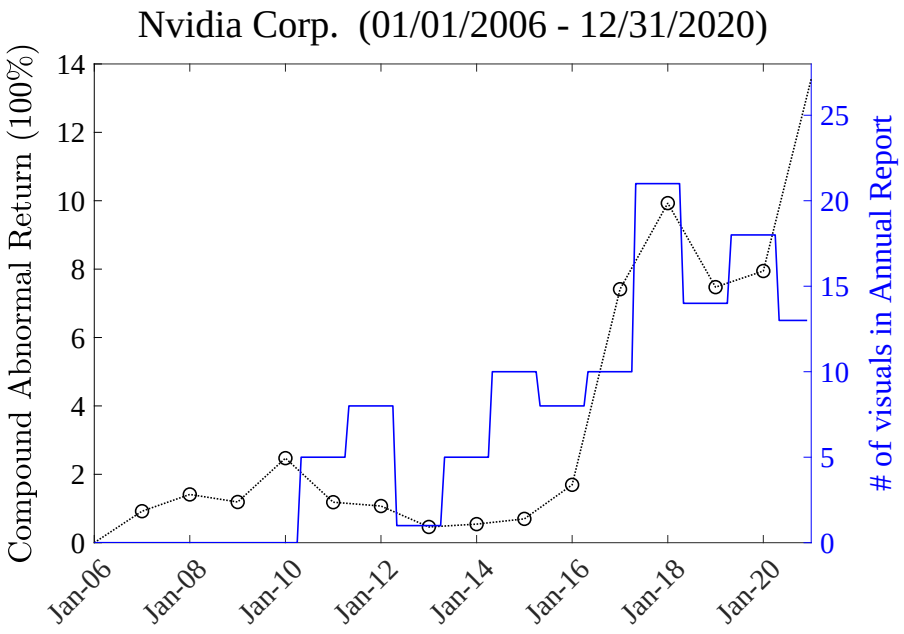


Figure 1. The compound abnormal returns of Nvidia's stock (dashed line) versus the numbers of visuals embedded in its annual reports (solid line) from 2006 to 2020.

fective disclosures, leading to reduced investment and potential stock undervaluation. This, in turn, creates a managerial incentive to adopt more engaging communication strategies, driving a substantial shift from textual to visual presentations. Not all types of data or investment decisions can be automated. Given limited cognitive resources (Kahneman, 1973; Sims, 2003), human investors tend to focus on data they can readily process while overlooking the information that is too complex to consume. As a result, information that receives little attention can become particularly valuable. Intuitively, fundamental analysis is more rewarding when the underlying information is more challenging for investors to interpret. This aligns with the observed strong yet delayed market responses and is further supported by our finding that positive abnormal returns are more pronounced for firms incorporating visuals that emphasize technology-related domains.

Our model captures the interplay between investors' inattention and firms' aesthetic communication strategies. As new efforts to engage investors, firms have even used videos and other creative elements in communications.[11] The positive abnormal returns indicate that such changes are effective for certain investors, though there remains significant room for improvement. Our findings may provide valuable insights for regulators when they consider

[11] As another anecdotal example, the German chemical company Covestro AG published a digital version of their 2019 annual report which incorporates video messages, interactive graphics, and even some quiz.

reforms to the current corporate disclosure framework. In addition, our study highlights that the shift from text-based to visually enhanced reporting is not merely decorative but represents a worthwhile improvement in corporate communication.

The past decade has witnessed a trend of applying content analysis in finance.[12] Most studies use textual analysis of firm disclosures or news media; see the survey by Loughran and McDonald (2016).[13] In contrast, non-textual analysis in financial research seems quite limited and is emerging most recently. Obaid and Pukthuanthong (2021) analyze the sentiment in news photos and find that their photo pessimism index predicts market return reversals and trading volumes, consistent with the behavioral theory (for example, De Long, Shleifer, Summers, and Waldmann (1990)). Edmans, Fernandez-Perez, Garel, and Indriawan (2022) propose a music-based sentiment measure and document relationships between music sentiment and stock returns. Jiang, Kelly, and Xiu (2021) apply image analysis methods on stock price charts to extract predictors for stock trends. In the real estate research, Carrillo (2008) and Benefield, Cain, and Johnson (2011) find that visuals have a large and positive effect on marketing outcomes. Adding a virtual tour may increase the transaction price and reduce the expected duration on the market.

Our work is contemporaneous but fills a void in the literature. Unlike music songs or news photos, we study graphics and visuals in firm annual reports which contain the most relevant and important corporate information. By law, the graphics and visuals embedded in those reports should be fair and accurate. In our sample, most graphics and visuals appear positive or at least neutral. This feature saves us from the technical problem of assigning positive or negative values to a large number of image files. To the best of our knowledge, we are the first study that provides a theoretical framework (see Kelly and Xiu (2023)) and uses unstructured data to understand the role of visual elements in financial statements. Following our study, there are only a few formal applications on visual information embedded in firm disclosures. For example, Cao, Cheng, Wang, Xia, and Yang (2024) analyze visual information in corporate executive presentations and responses from the stock markets. Ben-Rephael, Ronen, Ronen, and Zhou (2023) investigates the visual readability of annual reports and their impacts on information assimilation.

In the recent accounting literature, Nekrasov, Teoh, and Wu (2021) find that visuals in firms' Twitter earnings announcements can trigger more retweets and Google search queries.

[12]Content analysis can be applied to a wide range of "texts", including but not limited to written texts, images, animations, speeches, gestures, and music. There are two major branches. One is to analyze the verbal information. The other is to analyze the meaning derived from the use of multiple semiotic resources such as signs, emoticons, brand logos, photos, and other visual components (Ventola and Guijarro, 2009).

[13]These studies, for example, include Tetlock (2007), Tetlock, Saar-Tsechansky, and Macskassy (2008), Hanley and Hoberg (2010), Jiang, Lee, Martin, and Zhou (2019), and Huang, Tan, and Wermers (2020).

Christensen, Fronk, Lee, and Nelson (2021) document descriptive evidence on the trend of data visualization in 10-Ks. By Regulation S-T Rule 304, a typical 10-K does not contain large, colorful images. As a result, the sample size of graphics in 10-Ks is much smaller than that in firms' annual reports. Our findings suggest the necessity to carefully distinguish firms' visual annual reports from their plain 10-Ks. These two versions are not as interchangeable as it is widely perceived. The visual reports can be *informationally* larger than the 10-Ks.

Finally, our work contributes to the emerging literature on financial big data research (e.g., Begenau, Farboodi, and Veldkamp (2018), Martin and Nagel (2022)). Our sample has the features of big data as discussed by Goldstein, Spatt, and Ye (2021). Visual data is diverse, unstructured, and high-dimensional, with lots of interacting variables as one can envision. These properties pose challenges when researchers attempt to extract useful information from a large sample of visual data. In this regard, we take a step forward by demonstrating the value and predictive power of such a dataset, with both statistical and economic significance.

The rest of this paper is organized as follows. Section I discusses the institutional background. Section II presents our model and its predictions. Section III describes our data and methodology. Section IV presents our empirical results. Section V concludes. An Internet Appendix contains all the supporting information.

I. Institutional Background

Publicly listed companies must submit various reports (e.g., 10-Ks), schedules, and other filings electronically to the SEC, in compliance with Regulation S-T which outlines the rules and procedures pertaining to the SEC's EDGAR (Electronic Data Gathering, Analysis and Retrieval) system. Regulation S-T Rule 304 provides the regulatory code regarding the use of graphic, audio or video elements in corporate disclosures. To the best of our knowledge, this rule has not been mentioned in the academic literature. In Appendix A, we quote and discuss Rule 304, which has six key implications for our study, as summarized below:

1. **Equivalence in Principle.** Rule 304(a) acknowledges the use of visual content in firms' documents delivered to investors. However, there should be no material differences between such documents (e.g., graphic annual reports) and the corresponding electronic filings with the SEC (e.g., plain 10-Ks). By Rule 304(a), if the delivered version (of annual report) contains graphic content that is not reproduced in the electronic

filing (of 10-K), then the latter must simultaneously "*include a fair and accurate narrative description, tabular representation or transcript of the omitted material.*" From the legal perspective, the two documents are supposed to be informationally equivalent.

2. **Legal Liability.** Under Rule 304(b), the visual content in delivered annual reports "*is deemed part of the electronic filing and subject to the civil liability and anti-fraud provisions of the federal securities laws.*" This requires the visual content to be fair and accurate. Also, the narrative description, tabular representations or transcripts of the omitted material included in an electronic filing are deemed part of the filing.

3. **Subtle Differences.** Rule 304(b) states that "*However, to the extent such descriptions, representations or transcripts represent a good faith effort to fairly and accurately describe omitted graphic, image, audio or video material, they are not subject to the civil liability and anti-fraud provisions of the federal securities laws.*" The acknowledgement of "a good faith effort" may indicate the reality that both regulators and firms have to accept, considering that *one picture is worth a thousand words*. This item seems to leave room for some subtle information which may be embedded in the graphic annual reports but not readily visible in the plain 10-Ks.

4. **Easy Ignorance**. Under Rule 304(a), market participants may take for granted the "interchangeability" of visual annual reports and plain 10-Ks, and thus overlook the subtle differences between these two versions of disclosures. Under Rule 304(b), various subtle differences may be tolerated by regulators or may simply escape from their eyes.

5. **Record Keeping.** By Rule 304(c), firms are required to keep records for five years of their publicly distributed annual reports and make those accessible. This explains the abundance and continuity of our sample data.

6. **Plainness of 10-Ks**. Rules 304(d)(e)(f) set more stringent rules and specific guidelines for the electronic filings. For example, firms "*may not include animated graphics in any EDGAR document*". Rule 304(e) states that "*filers may not present in a graphic or image file information such as text or tables that users must be able to search and/or download into spreadsheet form (e.g., financial statements).*" These rules can push firms to follow the plain format in electronic filings, with minimal use of visuals.

Overall, Rule 304 gives much room for firms to use visuals and graphics in their published documents, provided that there is no material difference between those documents and their counterparts of electronic filings. However, there could be subtle differences between these two versions of disclosures. Rule 304, on the other hand, discourages the use of visuals and

graphics in electronic filings, consistent with the observed plain style of 10-Ks. The above institutional background explains why the sample data of visuals in our study is both large and informational.

II. Model

In this section, we develop a simple model that generates three key empirical predictions. We utilize a two-period framework of Kyle (1985) with multiple informed traders and introduce another player, the firm manager, who has the correct belief about the fundamental value of her firm's stock. One component of this value has not been recognized by the market yet, probably due to the intricate process of realizing its payoff or value. This component is subtle and not easily discernible. The manager cannot effectively convey this information through standard disclosures. However, through some other venue or channel, the manager can make extra efforts that help attentive traders to capture and capitalize this information.

Setup. Consider a market with one risky asset (stock) and two trading rounds, $t = 1, 2$. The stock fundamental value, $\tilde{v}$, is a sum of two independent Gaussian random variables, $\tilde{v} = \tilde{v}_0 + \tilde{\alpha}$. The first component, $\tilde{v}_0 \sim \mathcal{N}(1, \sigma_v^2)$, has been commonly recognized by the market and privately observed at time $t = 0$ by M attentive traders, indexed by $i = 1, ..., M$. The second component, $\tilde{\alpha} \sim \mathcal{N}(\overline{\alpha}, \sigma_\alpha^2)$, is subtle and not recognized by the market.

As a business insider, the manager has a clear vision of her firm, knowing that the stock value $\tilde{v}$ is actually drawn from the normal distribution $\mathcal{N}(1+\overline{\alpha}, \sigma_v^2+\sigma_\alpha^2)$. The manager does not observe the value of $\tilde{v}$ or trade this stock. She is aware that $\tilde{\alpha}$ has not been recognized by the market and the stock is thus mispriced by an average amount of $\overline{\alpha}$. The manager is eager to improve the market valuation of her firm. However, she may struggle to convey her belief of $\tilde{\alpha}$ to investors via plain disclosures. In an alternative venue, the manager can make extra efforts with indirect informational effects, in terms of catching investors' attention and promoting their research about $\tilde{\alpha}$. At $t = 0$, the manager makes a choice $s \in \{0, 1\}$ which is publicly observable. Choosing $s = 1$ incurs some cost $c > 0$ for the manager, but it may benefit traders who pay attention to the manager's action. After research, these traders may discover the value of alpha. Choosing $s = 0$ is costless, yet it raises no trader's interest and has no informational impact on any trader.[14] In our empirical context, $s = 1$ represents the manager's substantial use of enhanced visuals in annual report, while $s = 0$ means her

[14]This choice may withhold important proprietary information (Darrough and Stoughton (1990)) or avoid miscommunication with investors when firms are uncertain about their preferences (Bond and Zeng (2022)).

	$t=0$	$t=1$	$t=2$	$t=3$
Firm Manager	choose $s \in \{0,1\}$	...	...	see $v = v_0 + \alpha$
Attentive Traders	see v_0 and s	trade $x_{i,1}$	may see α; trade $x_{i,2}$	receive payoff
Market Makers	set $p_0 = 1$	set $p_1(y_1)$	set $p_2(y_1, y_2)$	see $v = v_0 + \alpha$

reduced use of visuals. Similar to Verrecchia (1983), the manager's objective function is

$$\max_{s\in\{0,1\}} \mathrm{E}^{\mathcal{F}}[\tilde{p}_2(s) - p_0] - s \cdot c, \tag{1}$$

where $\mathcal{F}$ represents the manager's correct belief that $\tilde{v}$ is drawn from $\mathcal{N}(1+\overline{\alpha}, \sigma_v^2 + \sigma_\alpha^2)$ and where $\tilde{p}_2(s)$ denotes the firm's stock price at $t=2$ indirectly affected by the choice of s.

The above table illustrates the players and the timeline of events in our model. While the manager's choice $s \in \{0,1\}$ is publicly observable, it is only noticed by attentive traders who have privately observed $\tilde{v}_0$ at the beginning. When $s=0$, each trader receives no hint and cannot find new information. When $s=1$, each trader may discover the value of alpha with probability $\mu > 0$, by research for one period of time. One may interpret the length of the first period as the minimum amount of time for traders to discover alpha.[15] Due to the subtlety of alpha, each trader may fail to find it with probability $1-\mu$. Each attentive trader can submit two market orders, $x_{i,1}$ and $x_{i,2}$, to maximize his expected trading profit, $\mathrm{E}^{\mathcal{A}}\left[(\tilde{v} - \tilde{p}_2) \cdot x_{i,2} + (\tilde{v} - \tilde{p}_1) \cdot x_{i,1}\right]$, where $\mathcal{A}$ represents the trader's belief system. Specifically, each trader observes both v_0 and s at $t=0$ and believes that v_0 is the actual stock value before his trading at $t=1$. Each trader observes v_0, s, p_1, and $s\alpha\mathbf{1}_\mu$ before trading at $t=2$, where $\mathbf{1}_\mu$ is a Bernoulli random variable that equals one with probability μ. Seeing $s=1$ does not necessarily change a trader's prior; his belief is updated only after he discovers alpha, as summarized by the product term $s\alpha\mathbf{1}_\mu$. The stock value, $\tilde{v} = \tilde{v}_0 + \tilde{\alpha}$, is realized and publicly revealed at $t=3$. In our model, attentive traders are fundamental investors in the sense that their trading decisions must be backed by material information (i.e., knowing the value of v_0 or α). The manager's action $s=1$ may give some useful hints for alpha.

At each time t, competitive market makers absorb the aggregate order flow, $\tilde{y}_t = \sum_{i=1}^{M} \tilde{x}_{i,t} + \tilde{u}_t$, where $\tilde{u}_1 \sim \mathcal{N}(0, \sigma_u^2)$ and $\tilde{u}_2 \sim \mathcal{N}(0, \sigma_u^2)$ are noise demands. All the random variables,

[15] In practice, delays and failures can occur throughout the research process and within the hierarchical structures governing institutional investors' decision-making. This aligns with theoretical insights from Radner (1993), Bolton and Dewatripont (1994), and Stein (2002). As documented by Evans et al. (2022), this hierarchy effect is more pronounced for fundamental investors than for quantitative traders.

$\tilde{v}_0$, $\tilde{\alpha}$, $\tilde{u}_1$, and $\tilde{u}_2$, are mutually independent, with σ_v and σ_u being common knowledge. The parameter μ is known among attentive traders and the manager. The initial stock price is $p_0 = 1$, based on the common prior $\tilde{v} \sim \mathcal{N}(1, \sigma_v^2)$. Without knowing the presence of alpha, market makers set $\tilde{p}_t$ to be their posterior expectation of $\tilde{v}$ conditional on order flows up to time t, that is, $\tilde{p}_1(\tilde{y}_1) = \mathrm{E}^{\mathcal{P}}[\tilde{v}|\tilde{y}_1]$ and $\tilde{p}_2(\tilde{y}_1, \tilde{y}_2) = \mathrm{E}^{\mathcal{P}}[\tilde{v}|\tilde{y}_1, \tilde{y}_2]$. Here, $\mathcal{P}$ denotes their common belief that $\tilde{v}$ is drawn from $\mathcal{N}(1, \sigma_v^2)$ and privately observed by M traders. The cumulative abnormal returns over the time window $[0, t]$ can be measured by $\mathrm{E}[(\tilde{p}_t - p_0)/p_0] = \mathrm{E}[\tilde{p}_t - 1]$.

Equilibrium. In our model, if a trader has found the alpha, he is not sure how many other traders have also found it. Each trader uses rational expectations to adjust his trading plan.

Proposition 1. *There exists a subgame perfect equilibrium where market makers set prices,*

$$\tilde{p}_1 = p_0 + \lambda_1 \tilde{y}_1, \qquad \tilde{p}_2 = \tilde{p}_1 + \lambda_2 \tilde{y}_2. \tag{2}$$

The price impact parameters, λ_1 and λ_2, are endogenously determined by M, σ_u, and σ_v. Each attentive trader adopts a pair of linear trading strategies $(X_{i,1}, X_{i,2})$, given by

$$X_{i,1} = \frac{\lambda_2(M+1)^2 - 2\lambda_1}{\lambda_2(M+1)^3 - 2M\lambda_1} \cdot \frac{\tilde{v}_0 - p_0}{\lambda_1}, \qquad X_{i,2} = \frac{\tilde{v}_0 - \tilde{p}_1}{\lambda_2(M+1)} + \frac{s\tilde{\alpha}\mathbf{1}_\mu}{\lambda_2[2 + \mu(M-1)]}, \tag{3}$$

where the product term, $s\tilde{\alpha}\mathbf{1}_\mu$, reflects the impact of manager's efforts $s = 1$ on each trader who may independently discover $\tilde{\alpha}$ with probability μ. Here, $\mathbf{1}_\mu$ denotes a Bernoulli random variable which equals 1 with probability μ and equals 0 with probability $1 - \mu$.

Proof. See Appendix H. □

Attentive traders may conduct fundamental analysis. Their trading decisions are executed only after confirming such decisions are backed by material information. This explains why in Equation (3) the initial trading strategy $X_{i,1}$ is only driven by $\tilde{v}_0 - p_0$ and independent of s but the next-period trading strategy $X_{i,2}$ has an extra term proportional to $s\tilde{\alpha}\mathbf{1}_\mu$. This term reflects a trader's potential knowledge of alpha. In our rational-expectations equilibrium, the manager understands the fundamental style of attentive traders. As a result, there will be no strategic or manipulative considerations when the manager chooses s.

Proposition 2. *Anticipating the trading equilibrium in Proposition 1, the manager chooses $s = \{0, 1\}$ based on the prior mean of alpha ($\overline{\alpha}$) relative to the threshold $K = \frac{2+\mu(M-1)}{\mu M} c > 0$.*

If $\overline{\alpha} \geq K$, the manager chooses $s = 1$ and the stock exhibits abnormal returns at $t = 2$:

$$\mathrm{E}[\tilde{p}_1 - p_0 \mid s = 1] = 0, \qquad \mathrm{E}[\tilde{p}_2 - p_0 \mid s = 1] = \frac{\mu M \overline{\alpha}}{2 + \mu(M-1)} > 0. \tag{4}$$

If $\overline{\alpha} < K$, the manager chooses $s = 0$ and the stock shows no abnormal returns.

Proof. See Appendix I. □

Equation (4) indicates a time delay in observing the abnormal returns of this stock. For this prediction to hold, the manager's informational motivation in choosing s must not be anticipated or exploited by speculators in the market. Speculative trading tends to be much faster than fundamental investing, as speculators often skip time-consuming research or information verification. Their trades can trigger abnormal returns at an earlier stage. In practice, traders are not restricted from betting on s since it is publicly observable. In our model, attentive traders refrain from doing so due to their investment style or discipline. They are classified as fundamental investors who trade only on material information.

Proposition 3. *If a speculator recognizes the equilibrium in Propositions 1 and 2, then he will trade at $t = 1$ to avoid competition at $t = 2$ with fundamental traders who remain uninformed about alpha when trading at $t = 1$. The speculator's trading demand has a positive lower bound, determined by his estimate of the manager's threshold K. Speculative trading can generate a positive abnormal return at an earlier stage after the manager chooses $s = 1$. If no abnormal return appears at $t = 1$, it implies that no speculator has recognized or exploited the signaling value of $s = 1$.*

In a rational-expectations equilibrium, agents may infer valuable information from equilibrium prices or other observable variables. In Proposition 3, the speculator is inferring a trading signal from the manager's action $s = 1$ which is observable and implicitly informational. Propositions 1, 2, and 3 imply a prediction: if the manager's informational motive in choosing $s = 1$ is not recognized or exploited by speculators, we may observe no abnormal returns in the earlier stage ($t = 1$) but positive abnormal returns in the later stage ($t = 2$).

In our model, investors can outperform the market using a fundamental strategy, which requires time-consuming research aided by hints from public data. One such source of hints may come from the manager's efforts ($s = 1$) to enhance firm communication. In our empirical study, the signal is extracted from the public data of enhanced visuals in annual reports. The profitability of this fundamental strategy may increase with the alpha's subtlety which

may reduce both μ and M and hence suppress traders' competition. Economically, M represents the overall attention from investors toward this firm's alpha, whereas μ reflects the effectiveness of firm communication with those investors. The probability μ can be influenced by various factors, including the manager's communication skills, firm reputation, regulatory environment (e.g., formatting requirement), and the organizational structure for investors' decision-making processes. In the extreme case of $\mu \to 0$, our model describes the situation of ineffective communication (e.g., poorly designed disclosures). When the perceived $\overline{\alpha}$ is high, a firm with $\mu \approx 0$ would be significantly undervalued by the market. This inefficient valuation can hurt firm performance so that the manager has incentive to improve the effectiveness (μ) of communication with investors. This then lowers the manager's threshold K for choosing $s = 1$, creating a noticeable shift from inefficient ($s = 0$) to efficient ($s = 1$) communication. In general, μ and M can be interrelated. A broader audience may incentivize greater managerial efforts to enhance communication effectiveness. In turn, more effective communication can attract more attentive investors, creating a reinforcing cycle between investor attention and firms' communication strategies. The cost parameter c can depend on σ_α^2 and μ. More volatile alpha may require more effective communication which needs more efforts. We can also model s as a continuous variable and allow c to increase with s. All those considerations do not change the main implications from our simple model.

Empirical Implications. We extend the model to a number of stocks. Each stock has its own alpha, denoted as $\tilde{\alpha}_n$ and indexed by $n = 1, ..., N$. The sample mean of alphas is zero. Each firm's alpha is acknowledged by its own manager but not by the market. Each manager chooses an action $s_n \in \{0, 1\}$ at the beginning based on her threshold, denoted K_n, which depends on parameters c_n, μ_n, and M_n. This model generates three testable predictions:

Prediction 1. *There are significant, positive abnormal returns for the group of firms whose managers have made efforts ($s = 1$) that provide certain hints for their firms' alphas.*

Managers are willing to make such efforts only if the expected benefits are substantial. Since the threshold is positive for each manager, the sample average abnormal return for those firms should be positive and statistically significant if the sample size is large enough. Empirically, the manager's action $s_n = 1$ is measured by the surge of visual elements in her firm's annual report. If this measure can predict stock returns, then we can tell that the manager's efforts have helped some investors to discover alpha.

Prediction 2. *For the group of firms whose managers have chosen to put no efforts ($s = 0$), their sample average abnormal returns are not statistically different from zero.*

Managers withhold information if the expected gains from their efforts are not promising. Their inaction raise no attention and leave no hints. As a result, the expected abnormal return is zero for each stock in this group. For a large sample, the abnormal return shall not be statistically different from zero. Empirically, the manager's inaction $s_n = 0$ is represented by the reduced use of visual elements in her firm's annual report. If we observe significant abnormal returns for firms with $s = 1$ (Prediction 1) and not for firms with $s = 0$ (Prediction 2), then we can confirm the critical role of enhanced visuals added by managers, given there is no pre-announcement drift in either test. This will also confirm our hypothesis of information asymmetry between firm managers (insiders) and financial markets (outsiders).

Prediction 3. *If the above return predictability of $s = 1$ is not expected or exploited by any speculators in this market and if the information about alpha is processed by fundamental investors, then there will be an apparent delay in observing the predicted abnormal returns. Moreover, the holdings of fundamental investors should increase for those firms with $s = 1$.*

The delayed emergence of anomaly is a key implication of Propositions 1, 2 and 3. The predicted increase in institutional holdings given $s = 1$ is an implication of Equation (3) in Proposition 1. If Prediction 3 is empirically confirmed, we can draw two critical implications. First, the delayed abnormal returns is likely caused by fundamental investors, since they do not immediately speculate on the hints inferred from the manager's action $s = 1$. Instead, they only trade on material information about alpha. Second, the return predictability of the manager's efforts is novel to the market since it has not been recognized by any speculators.

III. Data and Methodology

A. Data Collection of Annual Reports Visuals

We collect the printable versions of annual reports of 5,008 U.S. public firms over the period of 1994 to 2022. Most firms have such versions available as PDF files on their websites under the "Investor Relations" section. We search the Internet or request firms to send us their annual reports if not found on their websites.[16] We eventually have collected 47,715 annual reports which includes more than 6 million document pages.

[16]We complement the data by identifying annual reports on https://www.annualreports.com/, which covers more than 90% of firms listed on NYSE and NASDAQ.

To process roughly 6 million document pages, we customize a machine-learning algorithm which can automatically identify and analyze the graphics embedded in a PDF file. We mainly rely on the PyMuPDF and the Pillow library to process image contents.[17] PyMuPDF is a Python binding for MuPDF, a software package known for its top performance and high rendering quality. PyMuPDF also includes a built-in neural network based OCR engine. [18] PyMuPDF is instrumental in isolating the visual elements that are often key to financial insights. Pillow is an image library designed for access to data stored in a few basic pixel formats. It provides a general image processing tool. Our program first captures the graphic items in a file, including graphs, photos, and charts. Then it selects color-print items, filters out meaningless items such as titles and headings, calculates how much space the remaining graphics occupy on a page, and counts how many color-print graphics present in a document. Those output variables allow us to quantify the use of graphics in each PDF file. After analyzing more than 8 million image files, our program returns around 2 million image files that have high visual salience based on their colors and sizes. As a robustness check, we have also used a random subsample to manually confirm the accuracy of our computer algorithm.

After these images are extracted, they undergo a rigorous process of analysis. We deploy models developed by OpenAI, particularly the "CLIPProcessor" and "CLIPModel" from their CLIP (Contrastive Language–Image Pretraining) framework. These tools do more than just recognize image content; they intelligently enable us to categorize each image. This step is crucial as it allows us to organize the plethora of financial visuals into coherent categories, setting the stage for a more targeted analysis.

To deepen our understanding of the extracted images, we turn to Salesforce's BLIP (Bidirectional Language–Image Pretraining) model. This includes the "BlipProcessor" and "BlipForConditionalGeneration", which work together to generate concise, descriptive captions for each image. This enhancement is significant in qualitative analysis, making complex financial visuals more digestible and interpretable.

Our methodology also embraces an aesthetic evaluation of the images. This dimension of analysis begins with a simple heuristic assessment of brightness and contrast and then progresses to a more intricate aesthetic evaluation using an AI model (OpenAI). This model, specifically calibrated to assess aesthetic elements, offers a fresh lens through which to view the visual quality of financial reports – a perspective that traditional financial analysis rarely considers.

[17] See https://github.com/pymupdf/PyMuPDF and https://pypi.org/project/Pillow/.

[18] Optical Character Recognition (OCR) a digital technology that can convert different types of textual files, such as scanned paper documents, PDF files or digital images into editable and searchable data.

In our endeavor, each image is meticulously documented, capturing its dimensions, contents, aesthetic score, and more importantly, an AI-generated description. This information is compiled methodically into a CSV file, constructing a detailed database ripe for further analysis.

B. Image-level Content Analysis

With the process documented above, we finalize our initial sample with slightly more than 2 million image files with descriptive captions. To enhance data quality, we remove any captions that contain extraneous codes or characters without meaningful content, leaving us with 1,945,531 image files with unique and meaningful captions.

To systematically categorize these images, we leverage the capabilities of advanced natural language processing (NLP) models, specifically the ChatGPT-4o text learning model. Initially, we input all available captions into the NLP model to extract meaningful keywords that are prevalent throughout the dataset. Using these keywords, we develop a preliminary set of categories that capture the main themes and elements depicted in the images.

Subsequently, the model assists in refining these categories by identifying additional relevant keywords that align with the initial categorization but may not have emerged during the initial keyword extraction process. This iterative process ensures comprehensive coverage of possible content themes. For example, keywords such as "innovation", "spaceship", "satellite", "neural network", and "intelligence" were significant in identifying visuals related to the high-tech and innovation areas, while terms such as "medical", "mrna" and "antibiotics" were indicative of the healthcare and biotech areas. After defining these categories, we employ a text matching algorithm to further classify each image based on its caption, ensuring robustness in the categorization process and accounting for minor variations or errors in textual data.

Specifically, the categorization algorithm assigns images to thematic classes based on a set of logical conditions that detect the presence of indicative terms in each caption. For instance, captions that include terms such as "line graph," "bar chart," "timeline," or similar descriptors are classified as info-graphics. Similarly, terms that indicate nature scenes, urban landscapes, or physical structures lead to an assignment under the sceneries category. Categories such as people, transport, farming, and others are constructed similarly to capture the diverse visual content depicted in corporate annual reports. The final categorization includes 13 classifications based on more than 300 caption keywords. The detailed categories can be found in the Internet Appendix B.

In the following analysis, we can broadly categorize all the graphics and images in annual reports into two main groups:

1. Info-graphics: This group encompasses various forms of charts, figures, and exhibitions that visually present quantitative or qualitative information. Examples include bar charts, pie charts, histograms, timelines, and other illustrative tools designed to communicate data effectively.

2. Illustrative and Photographic Images (IP Visuals): This group includes all images that are not classified as info-graphics. These are primarily illustrative images, photos, and other photographic visuals depicting products, services, technologies, individuals, and relevant scenery. Examples include: Images of people, such as employees, customers, or leadership figures; visuals of landscapes or buildings, including headquarters, facilities, or operational environments; and representations of products, services, or technologies in action.

3. ITHB Visuals: This group is an important subset of IP visuals, and include visuals, photos, and representations of innovation, potential drug development and clinical trails, and other technologies in action. We also refer these images to R&D-focused visuals.

The ultimate categorization aims to uncover the types of visuals firms use to enhance the conveyance of financial and strategic information to stakeholders and shareholders, offering insight into how these images may affect investor perception and decision-making processes.

C. The Sample and Other Variables

We obtain from CRSP the monthly firm returns and market returns over the period of 1994 to 2022. We obtain the annual firm-level financial and accounting variables from Compustat, institutional holdings data from Thomson's CDA/Spectrum database (form 13F), Readability Indices are constructed using SEC Analytics and 10X summaries from Bill McDonald's website, patents data from USPTO and Kogan, Papanikolaou, Seru, and Stoffman (2017), and new product introductions data from Capital IQ. All the variables used in our empirical tests are defined in the Internet Appendix H. After excluding firms in the highly regulated sectors, such as financial (SIC codes 6000 to 6999) and utility (SIC codes 4901 to 4999) sectors, and firms without available earning announcement dates, our final sample contains 1,390 firms and 9,704 firm-year observations, which corresponds to roughly 278,000 image files studied. Table I describes our sample construction process.

[Insert Table I about here]

Table II presents the trends in the adoption of visual disclosure elements in annual reports by firms from 1994 to 2022, focusing on general visuals, info-graphics, and IP visuals. The dataset includes the total number of firms in the sample each year, as well as the breakdown of firms using specific types of visuals. Our sample includes only firms that ever produce annual reports, excluding firms that rely exclusively on regulatory filings such as 10-Ks. This sample selection may introduce upward bias in the adoption rates, as firms producing annual reports are more likely to employ enhanced visual disclosures.

[Insert Table II about here]

The sample size grew substantially, from 11 firms in 1994 to 702 firms in 2022. The sample coverage is limited prior to 2000. From 2000 to 2022, the proportion of firms using visuals has been around 70%, with a significant increase in the absolute number of firms adopting visual elements. This trend indicates that the use of visuals has become a standardized practice, with its growth leveling off as the sample size increased.

The adoption of info-graphics also grew, from 8 firms in 2000 to 283 firms in 2022. The growth in absolute numbers reflects increasing adoption, but the relatively low proportions, compared to IP images, suggest that info-graphics are less universally adopted than other visual elements.

The use of IP images exhibits a clear upward trajectory. In 2000, 132 firms (over 85%) employed IP images, and by 2022, this number rose to 487 firms (roughly 70%). Despite a slight proportional decline, IP images remained the most commonly used visual element across the sample, reflecting their importance in firms' strategic use of visual elements in their disclosures.

We also report the use of visuals across different sectors in the Internet Appendix .D. Consistent with the observations from Table II, we see the increasing reliance on visuals in annual reports across sectors. Notably, R&D-intensive sectors have exhibited rapid growth in visual adoption, possibly driven by their critical need to convey complex and nuanced information to stakeholders and shareholders.

Table III reports the summary statistics of the key variables for the sample we study. These include cumulative returns, FF5-factor alphas, the main measures S_{Images}, S_{Info}, S_{IP}, S_{ITHB}, *Visual Report*, firm characteristics (e.g., ROE, leverage, size, R&D expenses, firm age), readability of the annual reports/10-Ks, institutional attention and ownership, and proxies of innovation outcomes including the number of patents and new product introductions. For example, after the release of annual reports, the average 3-month and 6-month raw returns for all sample firms are found to be 5.7% and 9.5%. The FF5-factor alphas

based on weekly returns for the same event windows are 22.8 and 18.1 basis points. As documented earlier, most firm-year observations are associated with visual annual reports in our sample. However, only around 9.3% of observations correspond to the format-change events that firms newly issue visual reports.

[Insert Table III about here]

We also provide the logistic regression results in Appendix I, from which one can see that younger firms that have generally poorer 10-K readability from R&D intensive sectors are more likely to publish visually enhanced annual report. This is consistent with our theoretical intuition. Growth-oriented and younger firms are more likely to possess novel growth opportunities that are underappreciated by the market. In the context of our model, this implies that managers perceive a higher average alpha that remains unrecognized by investors. Additionally, firms in R&D-intensive sectors often deal with information that is novel, subtle, and nuanced—making it inherently more difficult for investors to interpret. This complexity tends to result in lower 10-K readability. Consequently, such firms may feel a stronger need to enhance communication with shareholders and stakeholders through the use of colorful and engaging visuals. In summary, characteristics such as strong growth, low readability, and an R&D-driven business model increase the likelihood that firms will rely on visuals in their annual reports, consistent with the predictions of our model.

IV. Empirical Results

A. Baseline Regression Results

We use panel regressions to formally test Prediction 1 and Prediction 2 from our model. To examine whether the graphics in annual reports contain information that predicts future stock returns, we estimate the following regression equation:

$$R_{i,t+\Delta} = \beta_0 + \beta_G \times S_{Images,i,t} + \beta_X \times X_{i,t} + \beta_F \times Firm_i + \beta_Y \times SIC2 \times Year_t + \epsilon_{i,t}, \quad (5)$$

where the dependent variable $R_{i,t+\Delta}$ is firm i's stock return in 3 months or 6 months after the report release date in year t. In this regression, the key independent variable is S_{Images}. As mentioned earlier, S_{Images} is an indicator variable that equals one for a firm-year observation if the firm newly issued visually enhanced annual reports in a fiscal year and that equals zero otherwise. The regression also includes several controlling covariates, denoted $X_{i,t}$, such as

firm size, return on equity, and leverage. To control for the textual effect on stock returns, we also include the Readability Index. All the regressions are estimated with both firm and industry × year fixed effects to mitigate the impact of unobserved firm-level heterogeneity and time-varying industry shocks, thereby reducing the omitted variable bias and isolating the causal impact of annual report enhanced visuals on future stock returns. The results are reported in Table IV. As shown in Columns (2) and (5), the coefficients β_G for S_{Images} are 1.815% and 0.091% in the six-month window. Both values are statistically significant, indicating that firms on average experience positive significant returns and abnormal returns (measured by Fama-French 5-factor Alpha) after they switch to issuing visual annual reports from using the plain 10-Ks. The regression coefficients of firm size and return on equity are also significant, indicating the effects of firm size and profitability. We do not observe an immediate stock price reaction to the release of newly visually enhanced annual report, reflecting that the new information is gradually incorporated into stock prices.

[Insert Table IV about here]

Collectively, the above empirical results show that the newly added visuals in firms' annual reports has predictive power for their future stock returns. By confirming both Prediction 1 and Prediction 2, these results confirm our theory that firms have intentionally employ extra visuals in their annual reports to help convey extra information which is not anticipated or noticed by the market. However, without identifying which types of visuals have the greatest predictive power for future returns, we cannot fully understand the underlying mechanism and managerial incentives for switching from plain 10-Ks to visually enhanced annual reports.

B. Categorizing Visuals - An Image Level Analysis

To understand mechanisms by which the inclusion of visuals in annual reports influences stock prices, we analyze the type, color, and contents of each image in every annual report. We closely examine each image, distinguishing between artistic elements (e.g., logos, team photos, cover pages, products, related product market sectors, scenery, animals, people, or services) and those designed for numerical and textual reinforcement (e.g., figures, tables, charts, and exhibitions). By ensuring that the images are sized to capture human attention, we decompose our primary measure, S_{Images}, into three components representing firms' efforts in releasing new visual content. These categories include info-graphics (S_{Info}); IP images (S_{IP}); and visuals focused on innovation, technology, and product advancements in IT, healthcare, and biotech (S_{ITHB}, a subset of IP images). Substituting the main inde-

pendent variable, we re-estimate baseline regressions using these measures, and the results are presented in Table V. Interestingly, the estimates of S_{Info} are not significantly different from zero, but both S_{IP} and S_{ITHB} are positive and statistically significant. The estimates of both S_{IP} and S_{ITHB} are also economically significant. The coefficient 2.205 of S_{IP} in column 2 indicates that firms adding IP visuals experience a 2.2% positive cumulative return on average, over the 6-month window after the release of new annual reports, and the estimate on S_{IP} of 0.102 shows the cumulative abnormal returns of adopting new IP visuals in annual reports as high as 2.69% $(=(1 + 0.102)^{26} - 1)$. $In particular, the estimates of \mathrm{S}_{ITHB}$ are even more economically pronounced, indicating cumulative returns almost at 5% and cumulative abnormal returns around 4%.

[Insert Table V about here]

Our findings indicate that info-graphics, which largely reiterate existing numerical data, do not predict returns, suggesting that firms are unlikely to use numerical tables and graphs to convey new insights about a firm's fundamentals. In contrast, professional images—especially those emphasizing innovation and technological advancements—are predictive of returns. This suggests that firms strategically use visual information as a channel to communicate subtle, nuanced details, such as insights related to R&D, that might otherwise be challenging for the public and investors to grasp.

C. Robustness

We conduct several robustness tests to verify the pattern of abnormal returns subsequent to the release of new visually enhanced annual report.

Matched Sample. To avoid sample selection bias and potential endogeneity issues stemming from unobserved firm and industry characteristics, we employ a matched sample as an alternative benchmark and re-estimate the regressions presented in Table V. Specifically, we compare firms that newly issue visually enhanced annual reports with matched industry peers that have similar characteristics but do not alter their reporting formats.[19] The results are reported in Panel A of Table VI. We re-estimate the same regressions as in Table V, incorporating firm and year fixed effects, given that the new subsample is matched within the same industry. Consistent with Table V, the regression coefficients for S_{Info} are not statistically significant, whereas the coefficients for both S_{IP} and S_{ITHB} are statistically significant in nearly all specifications, with the only exception of the FF-Alpha regression

[19]To ensure comparability, control firms are matched by industry, firm size, firm age, and 10-K readability in the year immediately preceding the treated firms' reporting format change.

on $S_{Pro-Images}$. Notably, the economic magnitude of S_{ITHB} in both specifications is almost double that of S_{IP}.

Subsample based on R&D Sectors. Given that our findings are more pronounced for visuals related to innovation and technology advancement in tech and biotech, a natural concern is that the results might be driven by R&D sectors, where firms tend to grow faster. To rule out this possibility, we restrict our testing sample to include only firms in R&D sectors.[20]. We then estimate the same regressions as in Table V based on this restricted subsample, and the results reported in Panel B of Table VI. Consistent with Table V, the regression coefficients for S_{Info} are not statistically significant, whereas the coefficients for both S_{IP} and S_{ITHB} are statistically significant in all specifications at the 5% or 1% levels. Moreover, the economic magnitude of S_{ITHB} is almost double that of S_{IP}. Collectively, these robustness tests confirm our findings from Table V and rule out alternative explanations.

[Insert Table VI about here]

D. Fundamental Investors

Our results demonstrate the predictive power of annual report enhanced visuals. This return predictability is not anticipated by the market. *Ex post*, it remains unknown who are actually influenced by the visually enhanced annual reports. Our model assumes that the extra information is identified and processed by fundamental investors, not by speculators or retail traders. In general, fundamental investors are highly concerned about the truthfulness and materiality of the information they possess. They want to confirm that their long-term investments are backed by superior, fundamental information. In contrast, speculators tend to bet on short-term opportunities, mostly based on their conjectures or inferences about those opportunities without solid evidence or information. If the abnormal returns identified in this paper had been recognized by some speculators, then their trading would have taken place shortly after companies initiated graphic reports, creating abnormal returns in an earlier stage. If, instead, the abnormal returns are novel to the market and only fundamental investors process the related information, then we expect an apparent delay in the emergence of abnormal returns, as stated in Prediction 3.

To test this prediction, we plot in Figure 2 the regression coefficients and corresponding

[20]We classify R&D sectors by including firms that meet the following SIC code criteria, aligning with the Fama-French 12 industries: firms are classified as R&D sectors if their SIC codes fall between 2830 and 2839, between 3570 and 3579, between 3660 and 3699, between 3810 and 3819, between 3841 and 3849, between 7370 and 7379, between 8070 and 8099, or if the code is exactly 7373.

confidence intervals for the variable S_{ITHB} from Table 5, spanning from the month firms initially release visually enhanced reports up to 12 months thereafter. As illustrated, the effect does not emerge immediately but after the second month; coefficients remain statistically significant for several subsequent months, indicating persistent abnormal returns. Consistent with Prediction 3, this result shows that there is no speculative trading on events of firms' reporting format changes and thus the return predictability of visually enhanced annual reports is indeed a novel finding to the general market. This pattern also aligns with the notion that visuals convey forward-looking information, as their predictive power gradually diminishes after reaching a peak.

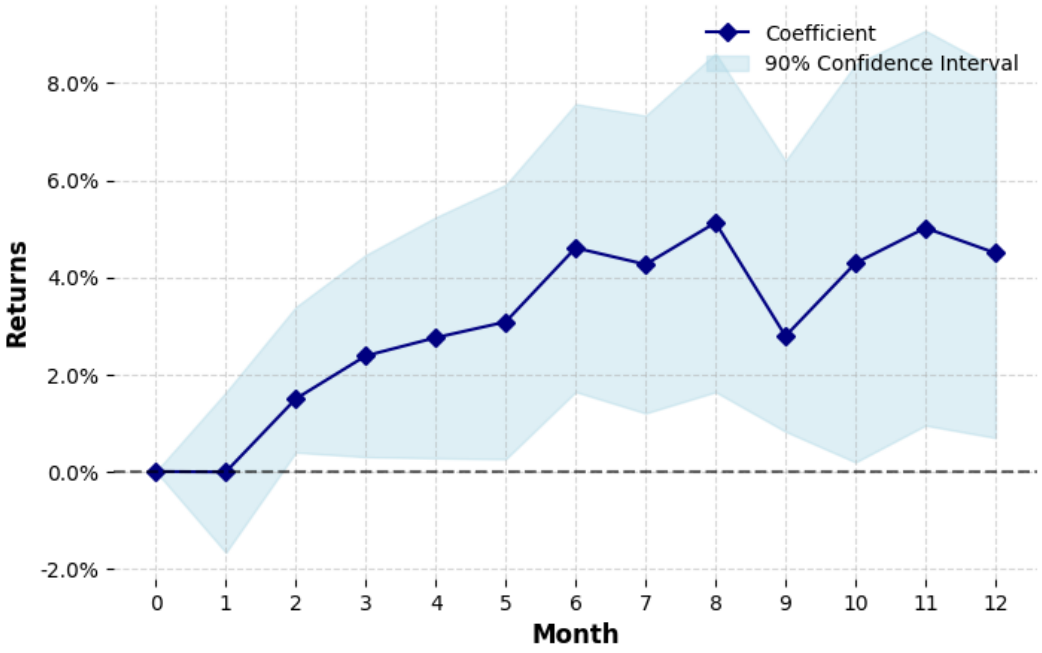


Figure 2. The estimated coefficients along with their corresponding 90% confidence intervals of the variable S_{ITHB} in the regression in Table V, across twelve time periods (from the month of annual report release to the 12th month after the release).

The delay in stock price response can be an important feature of fundamental investing. Fundamental investors may spend time collecting more data to substantiate their investment decisions. They may also spend extra time on analyzing soft information. According to Liberti and Petersen (2019), hard information includes quantitative data such as earnings numbers in financial statements. This type of information can be easily absorbed by traders. In contrast, it can be difficult to express (qualitative) soft information in numbers or in text. Comprehension of soft information also requires one's prior knowledge and even talents. As another friction, hierarchical organizations may also delay information communication

and investment decisions. Evans et al. (2022) show that the hierarchy effect is way more prominent for fundamental investors than for quantitative traders. This also agrees with theoretical implications of Radner (1993), Bolton and Dewatripont (1994), and Stein (2002).

The absence of speculative trading and the delay in stock price reactions suggest that the information driving the observed abnormal returns is most likely processed by fundamental investors. To examine this implication, we test whether annual report enhanced visuals, particularly those focused on R&D and technology advancement, leads to the post-event changes of fundamental investor behaviors such as institutional attentions and holdings.

We utilize the institutional attention data from Bloomberg and test if institutional investors do pay more attention to the focal firms who enhanced the visual usage in their annual reports. We estimate Poisson regressions for different types of visuals, and the dependent variable is the number of abnormal attention shocks during the given test window.[21] The results are reported in Table VII. The results demonstrate that illustrative and photographic images (S_{IP}) have a strong and highly significant positive association with abnormal institutional attention in both the first quarter (column 2, coefficient = 0.118, t=2.96) and the second quarter (column 5, coefficient = 0.091, t=2.63). This indicates institutional investors pay particular attention to IP visuals shortly after the release. Similar to our previous findings, R&D focused visuals (S_{ITHB}) attract even more investor attention (16.9% vs. 11.8% increase in attention) in both quarters (coefficients 0.169 in column 3 and 0.157 in column 6, both significant at 5%). Info-graphics (S_{Info}), however, do not show significant results, which is consistent with our return results.

[Insert Table VII about here]

Institutional investment decisions typically adjust more slowly than institutional attention. Therefore, we expect to observe a delayed effect on institutional holdings following an increase in their attention toward these firms. Given the quarterly data we can observe, we regress the percentage changes of firms' institutional ownership during one-quarter, two-quarters, and three-quarters after the release of new annual reports, together with other control variables similar to the regression equation (5). The results are reported in Table VIII. The coefficients of (S_{IP}) and (S_{ITHB}) is insignificant for the first quarter, and its magnitude keeps increasing with the length of post-event time window, and becomes statistically significant for both the second and the third quarter. Given that $Institutional\ Flow[3Q]$ has a mean of 0.147, the significant coefficient of S_{ITHB} 0.250 indicates that institutional

[21] We construct the institutional abnormal attention shocks following Ben-Rephael, Da, and Israelsen (2017). The details as to how to obtain the data from Bloomberg and variable construction can be found in Ben-Rephael et al. (2017).

investors increase their holdings in firms that include new R&D-focused visuals in their annual reports by more than double, relative to firms that do not. This pattern is consistent with Prediction 3 and with the observed abnormal returns in Table V.

[Insert Table VIII about here]

Combining the results in Tables V, VII, and VIII, we conclude that the observed abnormal returns (with a delay but no reversals) reflect the incorporation of nonpublic material information into stocks prices and this extra fundamental information is processed by fundamental investors (not speculators) who may receive hints or insights from visually enhanced reports.

E. Real Effects

If firms mainly use visually enhanced annual reports to hint at extra fundamental information, then the question for investors is whether those hints are credible. Specifically, do firms that release visual reports experience any improvements in their fundamentals subsequently? To answer this question, we estimate regressions where the dependent variables are the number of patents filed and eventually granted (measured at t+1 through t+3) and the number of new product introductions (measured at t+4 and t+5) in the upcoming years. Because innovation outcomes, such as patents, typically take several years from initial filing to eventual granting, and many patents filed in recent years remain pending, we conclude our analysis in 2020. Following the literature (e.g., see Cohn, Liu, and Wardlaw (2022)), we estimate the Poisson regressions as the number of patents are count-like variables. We control for firm-level fixed effects and SIC2 $\times$ year fixed effects, ensuring that unobserved heterogeneity do not drive the results. Table IX reports the results.

The coefficient of S_{ITHB} increases over time, becoming statistically significant starting at t+2 with a stronger significance at t+3 (t=2.32). The coefficient at t+3 of 0.155 can translate into a 16.8% increase in innovation outcome. It suggests that for firms from the same sector in the same year, the enhanced use of visuals related to R&D and technology advancement corresponds to a steady growth in innovation outcomes measured by the number of patents.

The results regarding new product introductions further reinforce this interpretation. Specifically, the coefficient of S_{ITHB} becomes positive and statistically significant at t+5 (coefficient = 0.146, t = 1.99), indicating that firms that adopt new R&D-focused visuals experience a significant increase in new product introductions. This coefficient corresponds to an 15.7% increase in the number of new product introductions. Given the nature that

most of new products are launched after firms have their patents approved, the delayed impact on new product introductions further supports our existing findings. Taken together, these results suggest a consistent pattern in which the enhanced use of R&D-focused visuals corresponds to tangible long-term innovation outcomes, both in patent productivity and new product introductions.

[Insert Table IX about here]

The above result is consistent with our theory. First, it provides more direct evidence that the visuals in the annual report are overall credible, as these visuals are directly related to innovation and technology advancements. It usually takes years for a new patentable technology to be filed and granted for patents, and eventually to be incorporated into new products. Therefore, the observed effect on patents and delayed effect on new products provides strong evidence that firms tend to use visuals to communicate subtle but important messages to shareholders and stakeholders.This can also explain why institutional investors may take into account such hints. The R&D operations are often complex and may involve proprietary information. Visuals and pictures seem suitable tools when firms intend to spread subtle messages about $R\&D$-related firm potentials.

There can be real benefits from using visual reports to signal $R\&D$-related fundamental improvements. This can reduce information asymmetry, boost stock prices, and lower costs of capital, which in turn promote firm investments (Fishman and Hagerty (1989), Diamond and Verrecchia (1991)). Firms may also reduce their perceived risk about investment opportunities and harness the feedback effects of financial markets (see Chen, Goldstein, and Jiang (2007), Bond, Edmans, and Goldstein (2012), Edmans, Goldstein, and Jiang (2012)).

F. Policy Implications

Our results support the managerial incentive to adopt a multimedia approach in preparing and presenting corporate information. Investors may have limited ability to filter valuable information from text-heavy reports (e.g., 10-Ks); they may also struggle with the prior knowledge required for understanding certain information. This may explain why the standard filings can be less attractive to investors, consistent with the finding by Loughran and McDonald (2017) that the activity of investors' requests for 10-Ks on EDGAR of the SEC is surprisingly low and also with the finding by Cohen, Malloy, and Nguyen (2020) that investors are inattentive to format changes in firms' financial reports.

Recognizing those problems, firms may invest more resources in producing attractive

reports. This practice shares similarities with marketing but has a different target audience. Visuals can help investors overcome their inattention and capture the big-picture insight. We provide the novel evidence on the positive economic effect of strategically using visuals in firms' annual reports. Appropriate use of visuals can enhance communication with investors and improve the market's informational efficiency. Hence, our findings may help regulators consider going beyond the current regulatory framework and accommodate the needs of firms to make more user-friendly presentations.[22]

Of course, there could be downside risk. The case of WeWork IPO filing is one example. The first Form S-1 that WeWork filed for its IPO is an image-heavy document, beyond the norm followed by most companies that go public. According to the news media,[23] the former WeWork executives spent months and a fortune producing the S-1 report "*like it was the September issue of Vogue.*" WeWork had hired a former director of photography at *Vanity Fair* magazine and sent photographers to reshoot the company's locations around the world. Its overuse of photos in the S-1 filing has been viewed by the financial media as glossy and inappropriate. Similar problems may also arise in the graphic annual reports of some firms.

Our empirical approach can expand regulators' tools that monitor misleading or abusive practices. As done in this paper, regulators may construct a database of graphics embedded in firms' disclosures. Regulators may also adopt the simple metric we defined to gauge the *Visual Reports* of firm documents in real time. This can help detect suspicious activities and urge firms to fairly use graphics in their disclosures.

V. Conclusion

Studies that employ content analysis of financial reports have largely focused on standard numerical and textual data, with advances made by Tetlock (2007) and others (see, e.g., Loughran and McDonald (2016)). This is likely due to the complexity of non-textual content (e.g., graphics) that has hindered extracting useful information from such type of dataset in the past. With recent advances in machine learning, our paper takes the challenge to analyze, for the first time, the visuals embedded in public firms' annual reports. We provide not only extensive empirical results, but also a novel theoretical model to explain why firms

[22] See https://www.sec.gov/divisions/corpfin/guidance/regs-tinterp.htm for the SEC Compliance and Disclosure Interpretations for Regulation S-T in 2015. Readers are also referred to *A Plain English Handbook: How to Create Clear SEC Disclosure Documents*; see https://www.sec.gov/reportspubs/investor-publications/newsextrahandbookhtm.html.

[23] See, for example, the media coverage and comments from the Bloomberg website at https://www.bloomberg.com/opinion/articles/2019-10-01/dorm-room-hedge-fund-had-its-own-values.

make the extra efforts to incorporate visuals in their reports to begin with.

We find that those firms who enhance the visual contents in annual reports earn significant and positive abnormal returns in the following one to two quarters. This finding is more pronounced for firms using visuals that focus on innovation and technology advancement, and is robust to various specifications. We also find that the positive abnormal returns are accompanied by an increase in holdings of institutional investors. Moreover, firms who increase the visuals have seen an increase in their research and development outcomes in the following years, reflecting the positive real effect of visually enhanced annual reports on corporate activities.

Our model helps explain the puzzle as to why a large fraction of firms continue to provide visually aesthetic but seemingly redundant financial reports to the general public. As suggested by our theory, firms are motivated in using visuals to convey information which is of value but not easily discernible, or to hint at potentials not reflected in their financial statements. Therefore, the added visuals in firms' annual reports are more than a decorative advertisement, but constitute a valuable source of information about firm fundamentals.

REFERENCES

Alldredge, Dallin M, and David C Cicero, 2015, Attentive insider trading, *Journal of Financial Economics* 115, 84–101.

Barber, Brad M, and Terrance Odean, 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *Review of Financial Studies* 21, 785–818.

Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp, 2018, Big data in finance and the growth of large firms, *Journal of Monetary Economics* 97, 71–87.

Ben-Rephael, Azi, Zhi Da, and Ryan D Israelsen, 2017, It depends on where you search: Institutional investor attention and underreaction to news, *The Review of Financial Studies* 30, 3009–3047.

Ben-Rephael, Azi, Joshua Ronen, Tavy Ronen, and Mi Zhou, 2023, Show me!" the informativeness of images, *Working Paper* .

Benefield, Justin D, Christopher L Cain, and Ken H Johnson, 2011, On the relationship between property price, time-on-market, and photo depictions in a multiple listing service, *Journal of Real Estate Finance and Economics* 43, 401–422.

Bolton, Patrick, and Mathias Dewatripont, 1994, The firm as a communication network, *Quarterly Journal of Economics* 109, 809–839.

Bond, Philip, Alex Edmans, and Itay Goldstein, 2012, The real effects of financial markets, *Annual Review of Financial Economics* 4, 339–360.

Bond, Philip, and Yao Zeng, 2022, Silence is safest: information disclosure when the audience's preferences are uncertain, *Journal of Financial Economics* 145, 178–193.

Cao, Sean S., Yichen Cheng, Meng Wang, Yusen Xia, and Baozhong Yang, 2024, Visual information and ai divide: Evidence from corporate executive presentations, *Available at SSRN* 4490834.

Carrillo, Paul E, 2008, Information and real estate transactions: the effects of pictures and virtual tours on home sales, *Department of Economics, The George Washington University, Washington, DC. Luettu* 14, 2016.

Chen, Qi, Itay Goldstein, and Wei Jiang, 2007, Price informativeness and investment sensitivity to stock price, *Review of Financial Studies* 20, 619–650.

Christensen, Theodore E, Karson Fronk, Joshua A Lee, and Karen K Nelson, 2021, Data visualization in 10-K filings, *Available at SSRN* 3748711.

Cohen, Lauren, Christopher Malloy, and Quoc Nguyen, 2020, Lazy prices, *Journal of Finance* 75, 1371–1415.

Cohn, Jonathan B., Zack Liu, and Malcolm I. Wardlaw, 2022, Count (and count-like) data in finance, *Journal of Financial Economics* 146, 529–551.

Da, Zhi, Joseph Engelberg, and Pengjie Gao, 2011, In search of attention, *Journal of Finance* 66, 1461–1499.

Darrough, Masako N, and Neal M Stoughton, 1990, Financial disclosure policy in an entry game, *Journal of Accounting and Economics* 12, 219–243.

De Long, J Bradford, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann, 1990, Noise trader risk in financial markets, *Journal of Political Economy* 98, 703–738.

Diamond, Douglas W, 1985, Optimal release of information by firms, *Journal of Finance* 40, 1071–1094.

Diamond, Douglas W, and Robert E Verrecchia, 1991, Disclosure, liquidity, and the cost of capital, *Journal of Finance* 46, 1325–1359.

Edmans, Alex, Adrian Fernandez-Perez, Alexandre Garel, and Ivan Indriawan, 2022, Music sentiment and stock returns around the world, *Journal of Financial Economics* 145, 234–254.

Edmans, Alex, Itay Goldstein, and Wei Jiang, 2012, The real effects of financial markets: The impact of prices on takeovers, *Journal of Finance* 67, 933–971.

Evans, Richard B, Martin Rohleder, Hendrik Tentesch, and Marco Wilkens, 2022, Diseconomies of scale, information processing and hierarchy costs: Evidence from asset management, *Journal of Financial and Quantitative Analysis* (Forthcoming).

Fishman, Michael J, and Kathleen M Hagerty, 1989, Disclosure decisions by firms and the competition for price efficiency, *Journal of Finance* 44, 633–646.

Friel, Susan N, Frances R Curcio, and George W Bright, 2001, Making sense of graphs: Critical factors influencing comprehension and instructional implications, *Journal for Research in Mathematics Education* 32, 124–158.

Goldstein, Itay, Chester S Spatt, and Mao Ye, 2021, Big data in finance, *Review of Financial Studies* 34, 3213–3225.

Grossman, Sanford J, 1981, The informational role of warranties and private disclosure about product quality, *Journal of Law and Economics* 24, 461–483.

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

Hawkins, David F, and Barbara A Hawkins, 1986, *The effectiveness of the annual report as a communication vehicle: A digest of the relevant literature* (Financial Executives Res Found).

Healey, Christopher, and James Enns, 2011, Attention and visual memory in visualization and computer graphics, *IEEE Transactions on Visualization and Computer Graphics* 18, 1170–1188.

Hirshleifer, David, and Siew Hong Teoh, 2003, Limited attention, information disclosure, and financial reporting, *Journal of Accounting and Economics* 36, 337–386.

Holden, Craig W, and Avanidhar Subrahmanyam, 1992, Long-lived private information and imperfect competition, *Journal of Finance* 47, 247–270.

Huang, Alan Guoming, Hongping Tan, and Russ Wermers, 2020, Institutional trading around corporate news: Evidence from textual analysis, *Review of Financial Studies* 33, 4627–4675.

Jiang, Fuwei, Joshua Lee, Xiumin Martin, and Guofu Zhou, 2019, Manager sentiment and stock returns, *Journal of Financial Economics* 132, 126–149.

Jiang, Jingwen, Bryan T Kelly, and Dacheng Xiu, 2021, (re-) imag (in) ing price trends, *Available at SSRN* 3756587.

Kahneman, Daniel, 1973, *Attention and effort*, volume 1063 (Prentice Hall, New Jersey).

Kastellec, Jonathan P, and Eduardo L Leoni, 2007, Using graphs instead of tables in political science, *Perspectives on Politics* 5, 755–771.

Kelly, Bryan, and Dacheng Xiu, 2023, Financial machine learning, *Foundations and Trends® in Finance* 13, 205–363.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *The Quarterly Journal of Economics* 132, 665–712.

Kosslyn, Stephen M, 1989, Understanding charts and graphs, *Applied Cognitive Psychology* 3, 185–225.

Kyle, Albert S, 1985, Continuous auctions and insider trading, *Econometrica* 1315–1335.

Larkin, Jill H, and Herbert A Simon, 1987, Why a diagram is (sometimes) worth ten thousand words, *Cognitive Science* 11, 65–100.

Liberti, José María, and Mitchell A Petersen, 2019, Information: Hard and soft, *Review of Corporate Finance Studies* 8, 1–41.

Lou, Dong, 2014, Attracting investor attention through advertising, *Review of Financial Studies* 27, 1797–1829.

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

Loughran, Tim, and Bill McDonald, 2017, The use of EDGAR filings by investors, *Journal of Behavioral Finance* 18, 231–248.

Martin, Ian WR, and Stefan Nagel, 2022, Market efficiency in the age of big data, *Journal of Financial Economics* 145, 154–177.

Milgrom, Paul R, 1981, Good news and bad news: Representation theorems and applications, *Bell Journal of Economics* 380–391.

Nekrasov, Alexander, Siew Hong Teoh, and Shijia Wu, 2021, Visuals and attention to earnings news on twitter, *Review of Accounting Studies* 1–43.

Obaid, Khaled, and Kuntara Pukthuanthong, 2021, A picture is worth a thousand words: Measuring investor sentiment by combining machine learning and photos from news, *Journal of Financial Economics* 144, 273–297.

Peng, Lin, and Wei Xiong, 2006, Investor attention, overconfidence and category learning, *Journal of Financial Economics* 80, 563–602.

Perilla, Juan R, and Schulten Klaus, 2017, Physical properties of the HIV-1 capsid from all-atom molecular dynamics simulations, *Nature Communications* 8, 1–10.

Pinker, Steven, 1990, A theory of graph comprehension, *Artificial Intelligence and the Future of Testing* 73–126.

Radner, Roy, 1993, The organization of decentralized information processing, *Econometrica* 1109–1146.

Roth, Wolff-Michael, and Gervase Michael Bowen, 2003, When are graphs worth ten thousand words? an expert-expert study, *Cognition and Instruction* 21, 429–473.

Schwabish, Jonathan A, 2014, An economist's guide to visualizing data, *Journal of Economic Perspectives* 28, 209–34.

Sims, Christopher A, 2003, Implications of rational inattention, *Journal of Monetary Economics* 50, 665–690.

Smiciklas, Mark, 2012, *The power of infographics: Using pictures to communicate and connect with your audiences* (Que Publishing).

Stein, Jeremy C, 2002, Information production and capital allocation: Decentralized versus hierarchical firms, *Journal of Finance* 57, 1891–1921.

Stenning, Keith, and Jon Oberlander, 1995, A cognitive theory of graphical and linguistic reasoning: Logic and implementation, *Cognitive Science* 19, 97–140.

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

Tetlock, Paul C, Maytal Saar-Tsechansky, and Sofus Macskassy, 2008, More than words: Quantifying language to measure firms' fundamentals, *Journal of Finance* 63, 1437–1467.

Tufte, Edward R, 1985, The visual display of quantitative information, *Journal for Healthcare Quality* 7, 15.

Ventola, Eija, and Arsenio Jesús Moya Guijarro, 2009, *The world told and the world shown: Multisemiotic issues* (Springer).

Verrecchia, Robert E, 1983, Discretionary disclosure, *Journal of Accounting and Economics* 5, 179–194.

Verrecchia, Robert E, 1990, Information quality and discretionary disclosure, *Journal of Accounting and Economics* 12, 365–380.

Verrecchia, Robert E, 2001, Essays on disclosure, *Journal of Accounting and Economics* 32, 97–180.

Table I: Sample Formation

This table presents the final sample construction over the period of 1994 to 2022, for both firm-year level information and image-level information.

# of image files with unique meaningful captions	1,945,531
# of categorized image files	1,151,630
# of categorized image files after excluding restrictive industries, ADRs, and stocks less than $5	818,233
# of categorized image files after excluding observations that we are unable to identify the annual report release dates	370,969
# of categorized image files, after excluding firms without accounting and return data, in the final sample	277,980
# of firms in the main analysis (firm-year observations)	1,390 (9,704)

Table II: The Use of Visuals by Year

This table presents the number of firms utilizing visuals for each year from 1994 to 2022. The second column shows the total number of firms in the sample for a given year, while the third through fifth columns display the total number of firms using visuals, info-graphics, and professional, colorful images or pictures, respectively.

Year	# of Firms	Using Visuals	Info-graphics	IP Images	ITHB Images
1994	11	6	1	6	3
1995	14	5	3	5	3
1996	29	14	5	14	8
1997	50	34	4	34	17
1998	73	59	5	59	34
1999	116	92	10	92	44
2000	168	132	8	132	82
2001	197	151	12	151	70
2002	241	193	16	192	100
2003	301	232	9	232	116
2004	333	256	17	256	129
2005	394	303	25	302	150
2006	450	336	40	336	135
2007	474	346	48	346	156
2008	467	331	49	330	147
2009	544	366	63	365	156
2010	555	391	94	390	171
2011	556	386	111	385	163
2012	569	406	154	404	151
2013	599	412	185	409	160
2014	631	445	221	442	165
2015	666	483	261	479	172
2016	638	461	240	459	170
2017	659	443	242	439	179
2018	650	435	235	429	165
2019	638	444	251	438	176
2020	668	440	229	434	210
2021	590	419	213	412	192
2022	702	499	283	487	235

Table III: Descriptive Statistics

This table reports the descriptive statistics of all major variables for all the sample firms we study. The definitions of variables are provided in the Internet Appendix H.

	N	Mean	Median	STD	P25	P75
Return [0,3]	9,548	5.701	5.322	19.13	-6.243	16.95
Return [0,6]	9,704	9.498	8.674	26.40	-6.64	24.61
Alpha [0,3]	9,330	0.228	0.180	1.63	-0.73	1.11
Alpha [0,6]	9,493	0.181	0.108	1.03	-0.45	0.74
S_{Images}	9,704	0.093	0.000	0.29	0.000	0.000
S_{Info}	9,704	0.028	0.000	0.17	0.000	0.000
S_{IP}	9,704	0.066	0.000	0.25	0.000	0.000
S_{ITHB}	9,704	0.021	0.000	0.14	0.000	0.000
Visual Report	9,704	0.711	1.000	0.45	0.000	1.000
# of Pages	9,704	5.72	3.00	6.94	1.00	8.00
Attention[1Q]	3,434	6.04	5.00	5.84	1.00	9.00
Attention[2Q]	3,441	11.80	9.00	11.16	2.00	18.00
Institutional Flow[2Q]	9,704	0.134	0.000	1.73	0.00	0.00
Institutional Flow[3Q]	9,704	0.147	0.000	2.04	0.00	0.00
Patents	7,258	51.86	2.00	182.2	0.00	18.00
New Products	7,723	1.01	0.00	1.80	0.00	1.00
Readability	9,704	17.26	16.86	5.24	13.49	20.65
Size	9,704	7.23	7.17	1.60	6.00	8.36
ROE	9,704	26.13	25.29	31.54	13.61	39.68
Leverage	9,704	30.56	28.66	25.69	5.03	47.22
R&D	9,704	5.54	1.22	8.72	0.00	7.52
Firm Age	9,704	23.15	19.00	17.04	9.00	35.00
INST	9,177	77.16	82.70	22.79	66.89	92.73

Table IV: Visually Enhanced Annual Report and Stock Return Predictability

This table reports the results of panel regressions that examine whether the annual report enhanced visuals predict future stock returns. Dependent variables are cumulative returns and 5-factor Alphas for the periods of 3-months, 6-months, and 9-months after the release of annual reports. All regressions are estimated with firm and SIC2 × year fixed effects. Standard errors are adjusted for heteroskedasticity and clustered by firm. *, **, *** indicate significance at the 10%, 5%, and 1% levels respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Return [0,3]	Return [0,6]	Return [0,9]	Alpha [0,3]	Alpha [0,6]	Alpha [0,9]
S_{Images}	0.568	1.815*	1.123	0.068	0.091**	0.039
	(0.79)	(1.87)	(0.97)	(1.04)	(2.25)	(1.16)
Size	-3.936***	-6.544***	-9.924***	-0.316***	-0.265***	-0.236***
	(-6.10)	(-7.74)	(-8.93)	(-5.99)	(-7.65)	(-8.39)
ROE	-0.029**	-0.028*	-0.025	-0.002*	-0.002***	-0.001**
	(-2.45)	(-1.87)	(-1.39)	(-1.71)	(-2.96)	(-2.54)
Leverage	0.050***	0.069***	0.101***	0.004**	0.003***	0.003***
	(2.68)	(2.78)	(3.50)	(2.55)	(3.46)	(3.32)
Readability	0.072	0.087	0.132*	0.005	0.005**	0.003
	(1.62)	(1.44)	(1.81)	(1.27)	(2.11)	(1.35)
# of Pages	-0.431	-1.330***	-1.622***	-0.057**	-0.035**	-0.028**
	(-1.30)	(-3.09)	(-3.40)	(-2.00)	(-2.02)	(-2.15)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
SIC2 x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	7,888	8,032	7,955	7,726	7,864	7,891
Adjusted R-Sq.	0.127	0.164	0.227	0.060	0.112	0.128

Table V: Categorizing Visuals

This table reports the panel regression results, using three measures quantifying the aggregated use of visuals and capturing the different categories of visuals. Dependent variables are 6-month cumulative stock returns and 5-factor weekly Alphas after the release of annual reports. Standard errors are adjusted for heteroskedasticity and clustered by firm. All regressions are estimated with firm and SIC2 × year fixed effects. *, **, *** indicate significance at the 10%, 5%, and 1% levels respectively.

	(1) Return [0,6]	(2) Return [0,6]	(3) Return [0,6]	(4) Alpha [0,6]	(5) Alpha [0,6]	(6) Alpha [0,6]
S_{Info}	1.118			0.103		
	(0.67)			(1.51)		
S_{IP}		2.205**			0.102**	
		(1.99)			(2.24)	
S_{ITHB}			4.612**			0.131*
			(2.56)			(1.68)
Size	-6.533***	-6.535***	-6.534***	-0.265***	-0.265***	-0.264***
	(-7.75)	(-7.74)	(-7.75)	(-7.67)	(-7.64)	(-7.65)
ROE	-0.029*	-0.028*	-0.028*	-0.002***	-0.002***	-0.002***
	(-1.90)	(-1.84)	(-1.86)	(-3.00)	(-2.92)	(-2.97)
Leverage	0.069***	0.069***	0.068***	0.003***	0.003***	0.003***
	(2.80)	(2.78)	(2.75)	(3.46)	(3.46)	(3.46)
Readability	0.086	0.087	0.087	0.005**	0.005**	0.005**
	(1.43)	(1.44)	(1.43)	(2.10)	(2.10)	(2.09)
# of Pages	-1.280***	-1.343***	-1.372***	-0.034*	-0.036**	-0.035**
	(-2.98)	(-3.11)	(-3.14)	(-1.92)	(-2.04)	(-2.02)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
SIC2 x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	8,032	8,032	8,032	7,864	7,864	7,864
Adjusted R-Sq.	0.164	0.164	0.164	0.111	0.112	0.111

Table VI: Robustness Tests

This table reports the results of two robustness tests. Dependent variables are 6-month cumulative returns and 5-factor Alphas. Panel A reports the results based on a matched sample where a control firm, in the same industry and with similar size and 10-K readability, is matched to a treated firm, and Panel B reports the results based on a subsample of R&D intensive sectors. All regressions are estimated with firm and SIC2 × year fixed effects (only year FE for the matched-sample regression). Standard errors are adjusted for heteroskedasticity and clustered by firm. *, **, *** indicate significance at the 10%, 5%, and 1% levels respectively.

Panel A:Regression Results based on the Matched Sample

	(1) Return [0,6]	(2) Return [0,6]	(3) Return [0,6]	(4) Alpha [0,6]	(5) Alpha [0,6]	(6) Alpha [0,6]
S_{Info}	-0.219			0.034		
	(-0.10)			(0.36)		
S_{IP}		3.280**			0.053	
		(2.04)			(0.77)	
S_{ITHB}			6.083***			0.195*
			(2.70)			(1.80)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	1,265	1,265	1,265	1,239	1,239	1,239
Adjusted R-Sq.	0.144	0.148	0.150	0.040	0.041	0.044

Panel B:Regression Results within R&D Intensive Sectors

	(1) Return [0,6]	(2) Return [0,6]	(3) Return [0,6]	(4) Alpha [0,6]	(5) Alpha [0,6]	(6) Alpha [0,6]
S_{Info}	0.871			0.101		
	(0.32)			(0.96)		
S_{IP}		3.745**			0.211***	
		(2.02)			(2.84)	
S_{ITHB}			6.326**			0.352***
			(2.21)			(2.87)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
SIC2 x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	2,940	2,940	2,940	2,843	2,843	2,843
Adjusted R-Sq.	0.135	0.136	0.136	0.104	0.106	0.106

Table VII: Visually Enhanced Annual Report and Institutional Attention

This table reports the Poisson regression results on which categorized visuals are linked with abnormal institutional investors' attention. Dependent variables are the aggregated attention shock counts for the period of one quarter (1Q) and two quarters (2Q) after the release of annual reports. Standard errors are adjusted for heteroskedasticity and clustered by firm. Firm and SIC2 x year fixed effects are included. *, **, *** indicate significance at the 10%, 5%, and 1% levels respectively.

	(1) Attention[1Q]	(2) Attention[1Q]	(3) Attention[1Q]	(4) Attention[2Q]	(5) Attention[2Q]	(6) Attention[2Q]
S_{Info}	0.076			0.003		
	(1.28)			(0.07)		
S_{IP}		0.118***			0.091***	
		(2.96)			(2.63)	
S_{ITHB}			0.169**			0.157**
			(2.22)			(2.50)
Size	0.155***	0.159***	0.157***	0.120**	0.123**	0.121**
	(2.71)	(2.78)	(2.75)	(2.28)	(2.32)	(2.31)
ROE	-0.001*	-0.001**	-0.001*	-0.001*	-0.001*	-0.001
	(-1.91)	(-1.98)	(-1.83)	(-1.70)	(-1.76)	(-1.63)
Leverage	0.002*	0.003*	0.003**	0.002*	0.002**	0.002**
	(1.87)	(1.93)	(2.00)	(1.93)	(1.96)	(2.04)
Readability	-0.001	-0.001	-0.001	-0.001	-0.002	-0.002
	(-0.38)	(-0.43)	(-0.42)	(-0.77)	(-0.82)	(-0.80)
INST	-0.000	-0.000	-0.000	0.000	0.000	0.000
	(-0.28)	(-0.30)	(-0.25)	(0.38)	(0.38)	(0.43)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
SIC2 x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	2,822	2,822	2,822	2,906	2,906	2,906
Pseudo R-Sq.	0.462	0.463	0.463	0.594	0.594	0.594

Table VIII: Visually Enhanced Annual Report and Institutional Investing

This table reports the panel regression results that examine whether firms' efforts of adding visuals are linked to the increase in fundamental investment inflow. Dependent variables are the *Institutional Flow* for the period of one quarter (1Q), two quarters (2Q), and three quarters (3Q) after the release of annual reports. Standard errors are adjusted for heteroskedasticity and clustered by firm. Firm and SIC2 x year fixed effects are included. *, **, *** indicate significance at the 10%, 5%, and 1% levels respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Inst Flow[1Q]	Inst Flow[1Q]	Inst Flow[2Q]	Inst Flow[2Q]	Inst Flow[3Q]	Inst Flow[3Q]
S_{IP}	0.001		0.171**		0.169**	
	(0.01)		(2.45)		(2.24)	
S_{ITHB}		0.081		0.224**		0.250**
		(1.04)		(2.10)		(2.09)
Size	-0.120***	-0.120***	-0.124**	-0.125**	-0.131**	-0.131**
	(-3.11)	(-3.12)	(-2.32)	(-2.32)	(-2.18)	(-2.19)
ROE	-0.000	-0.000	-0.000	-0.001	-0.002	-0.002
	(-0.38)	(-0.37)	(-0.43)	(-0.45)	(-1.34)	(-1.35)
Leverage	0.002*	0.002	0.003*	0.003*	0.004*	0.004*
	(1.65)	(1.63)	(1.83)	(1.82)	(1.87)	(1.86)
Readability	-0.000	-0.000	-0.002	-0.002	0.002	0.002
	(-0.09)	(-0.09)	(-0.54)	(-0.53)	(0.46)	(0.47)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
SIC2 x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	9,259	9,259	9,259	9,259	9,259	9,259
Adjusted R-Sq.	0.083	0.083	0.086	0.086	0.074	0.074

Table IX: The Real Effects of Annual Report Enhanced Visuals

This table reports the Poisson regression results that examine the relationship between annual report enhanced visuals and real corporate outcome variables measured by the number of patents filed that are eventually granted and the number of new product introductions. Table headings indicate dependent variables. Standard errors are adjusted for heteroskedasticity and clustered by firm. Firm and SIC2 x year fixed effects are included. *, **, *** indicate significance at the 10%, 5%, and 1% levels respectively.

	T+1 Patents	T+2 Patents	T+3 Patents	T+4 New Products	T+5 New Products
S_{ITHB}	-0.059	0.081*	0.155**	0.063	0.146**
	(-1.25)	(1.70)	(2.32)	(0.69)	(1.99)
Size	0.409***	0.313***	0.170**	0.011	-0.005
	(6.24)	(4.73)	(2.36)	(0.20)	(-0.10)
R&D	-0.006	-0.012	-0.023	-0.002	-0.004
	(-0.45)	(-0.89)	(-1.40)	(-0.35)	(-0.68)
ROE	0.000***	0.000***	0.000**	0.000	0.000
	(3.84)	(3.40)	(2.17)	(0.92)	(0.24)
Leverage	-0.000	-0.000	-0.001	0.000	-0.000
	(-1.25)	(-0.70)	(-0.76)	(1.05)	(-0.44)
Firm Age	-0.007	0.042	0.190	0.080	0.077
	(-0.05)	(0.31)	(1.42)	(0.92)	(0.91)
Firm FE	Yes	Yes	Yes	Yes	Yes
SIC2 x Year FE	Yes	Yes	Yes	Yes	Yes
N	5,502	5,005	4,590	4,485	3,979
Pseudo R-Sq.	0.949	0.949	0.950	0.407	0.409

Internet Appendix

Table of Contents

A Regulation S-T Rule 304 on the Use of Graphics and Images in Firm Disclosures 44

B The Categorization of the Keywords from Image Captions 47

C The Example of Categorized Images 48

D The Use of Visuals Across Different Sectors 54

E The Example of Nvidia's Annual Reports from 2010 to 2013 55

F Cover Page of Nvidia's Annual Report/10-K Issued in 2009 56

G Selected PRO-Images about Nvidia's Products and Services 57

H Selected Info-graphics about Performance and Compensations 58

I Proof of Proposition 1 59

J Proof of Proposition 2 62

K Definitions of Variables 63

L Supplementary Analysis 64

Appendix A. Regulation S-T Rule 304 on the Use of Graphics and Images in Firm Disclosures

Publicly listed companies are required to submit various filings to the SEC and comply with Regulation S-T which outlines rules and procedures pertaining to the SEC's EDGAR system. Among others, Rule 304 of Regulation S-T provides the regulatory details regarding the use of graphic, image, audio or video material in firm disclosures.[24]

Rule 304(a) states that "*If a filer includes graphic, image, audio or video material in a document delivered to investors and others that is not reproduced in an electronic filing, the electronically filed version of that document must include a fair and accurate narrative description, tabular representation or transcript of the omitted material. Such descriptions, representations or transcripts may be included in the text of the electronic filing at the point where the graphic, image, audio or video material is presented in the delivered version, or they may be listed in an appendix to the electronic filing. Immaterial differences between the delivered and electronically filed versions, such as pagination, color, type size or style, or corporate logo need not be described.*" As a note, "*if the omitted graphic, image, audio or video material includes data, filers must include a tabular representation or other appropriate representation of that data in the electronically filed version of the document.*"

Rule 304(a) acknowledges the use of graphics in firms' documents distributed to investors. However, there should be no material differences between those documents (e.g., user-friendly annual reports) and their corresponding electronic filings with the SEC (e.g., plain 10-Ks). Thus, if the delivered version of firm annual report contains graphic content that is not reproduced in the electronic 10-K filing, the latter must simultaneously include a fair and accurate narrative description, tabular representation or transcript of the omitted (graphic) content. From the legal perspective, these two documents should be informationally equivalent.

Rule 304(b) states that "*(1) The graphic, image, audio and video material in the version of a document delivered to investors and others is deemed part of the electronic filing and subject to the civil liability and anti-fraud provisions of the federal securities laws. (2) Narrative descriptions, tabular representations or transcripts of graphic, image, audio and video material included in an electronic filing or appendix thereto also are deemed part of the filing. However, to the extent such descriptions, representations or transcripts represent a good faith effort to fairly and accurately describe omitted graphic, image, audio or video material, they are not subject to the civil liability and anti-fraud provisions of the federal*

[24] See the Code of Federal Regulations for Regulation S-T Rule 304 at https://www.ecfr.gov/current/title-17/chapter-II/part-232/subject-group-ECFRd7d61c154ee5265/section-232.304.

securities laws." By Rule 304(b), firms are legally accountable for the graphic content in documents they deliver to investors. The SEC expects "a good faith effort" by firms to disclose the extra graphic content in their electronic filings. This seem leave room for subtle differences between the graphic annual report and the plain 10-K filing.

Rule 304(c) states that "*An electronic filer must retain for a period of five years a copy of each publicly distributed document, in the format used, that contains graphic, image, audio or video material where such material is not included in the version filed with the Commission. The five-year period shall commence as of the filing date, or the date that appears on the document, whichever is later. Upon request, an electronic filer shall furnish to the Commission or its staff a copy of any or all of the documents contained in the file.*" By Rule 304(c), firms must keep records for 5 years of their distributed annual reports and make those accessible. This explains the abundance and continuity of our sample data.

Rule 304(d) states that "*For electronically filed ASCII documents, the performance graph that is to appear in registrant annual reports to security holders required by Exchange Act Rule 14a-3 or Exchange Act Rule 14c-3 to precede or accompany proxy statements or information statements relating to annual meetings of security holders at which directors are to be elected (or special meetings or written consents in lieu of such meetings), as required by Item 201(e) of Regulation S-K, and the line graph that is to appear in registrant annual reports to security holders, as required by paragraph (b)(7)(ii) of Item 27 of Form N-1A, must be furnished to the Commission by presenting the data in tabular or chart form within the electronic ASCII document, in compliance with paragraph (a) of this section and the formatting requirements of the EDGAR Filer Manual.*"[25]

Rule 304(e) states that "*Notwithstanding the provisions of paragraphs (a) through (d) of this section, electronically filed HTML documents must present the following information in an HTML graphic or image file within the electronic submission in compliance with the formatting requirements of the EDGAR Filer Manual: The performance graph that is to appear in registrant annual reports to security holders required by Exchange Act Rule 14a-3 or Exchange Act Rule 14c-3 to precede or accompany registrant proxy statements or information statements relating to annual meetings of security holders at which directors are to be elected (or special meetings or written consents in lieu of such meetings), as required by Item 201(e) of Regulation S-K; the line graph that is to appear in registrant annual reports to security holders, as required by paragraph (b)(7)(ii) of Item 27 of Form N-1A; and any other graphic material required by rule or form to be filed with the Commission. Filers may, but are not required to, submit any other graphic material in a HTML document by presenting the*

[25] See https://www.sec.gov/edgar/filer-information/current-edgar-filer-manual.

data in an HTML graphic or image file within the electronic filing, in compliance with the formatting requirements of the EDGAR Filer Manual. However, filers may not present in a graphic or image file information such as text or tables that users must be able to search and/or download into spreadsheet form (e.g., financial statements); filers must present such material as text in an ASCII document or as text or an HTML table in an HTML document."

This paragraph (e) has been discussed and explained by the SEC.[26] In the SEC's view, "*'information such as text or tables that users must be able to search and/or download' consists of all information that the filer is required to include in the particular filing, such as disclosures in response to applicable form and Regulation S-K items and any additional information required to be included under Securities Act Rule 408 or Exchange Act Rule 12b-20.*" The SEC does "*recognize that registrants may present information in Commission filings in formats such as bar graphs or other graphics that are not text-searchable but that they believe may be more useful to readers than tables or paragraphs that present the same information in searchable form. We welcome the use of graphic and image files to make information more accessible for users, provided that the filing complies with applicable EDGAR size and formatting requirements and the filer's use of graphics does not interfere with a user's ability to search required information. Therefore, with regard to required disclosures, a filer may present required information using graphics that are not text-searchable and still comply with Rule 304(e) if the filer also presents the same information as searchable text or in a searchable table within the filing. The searchable information could be included, for example, together with the related graphic in the filing or in an appendix to the filing. Any additional information that the filer chooses to include in the filing and that is not required to be disclosed may be presented graphically without a separate text-searchable presentation.*" By 304(f), "*electronic filers may not include animated graphics in any EDGAR document.*"

In summary, Rules 304(a)(b)(c) describe the legal environment when firms use graphics in documents they delivered to investors. Those delivered documents (e.g., graphic annual reports) should be substantially equivalent to their counterparts of electronic filings (e.g., plain 10-Ks). Nonetheless, there could be subtle informational differences between these two which can be overlooked by investors and regulators. Rules 304(d)(e)(f) set specific and stringent rules for the use of graphics in electronic filings. In general, the graphic presentation of information is discouraged in those filings (e.g., 10-Ks). On one hand, the presentation of required information (e.g., financial statements) within electronic filings must be searchable or downloadable in a spreadsheet format. On the other hand, cautious firms may want to avoid litigation risk and minimize the use of graphics in their electronic filings with the SEC.

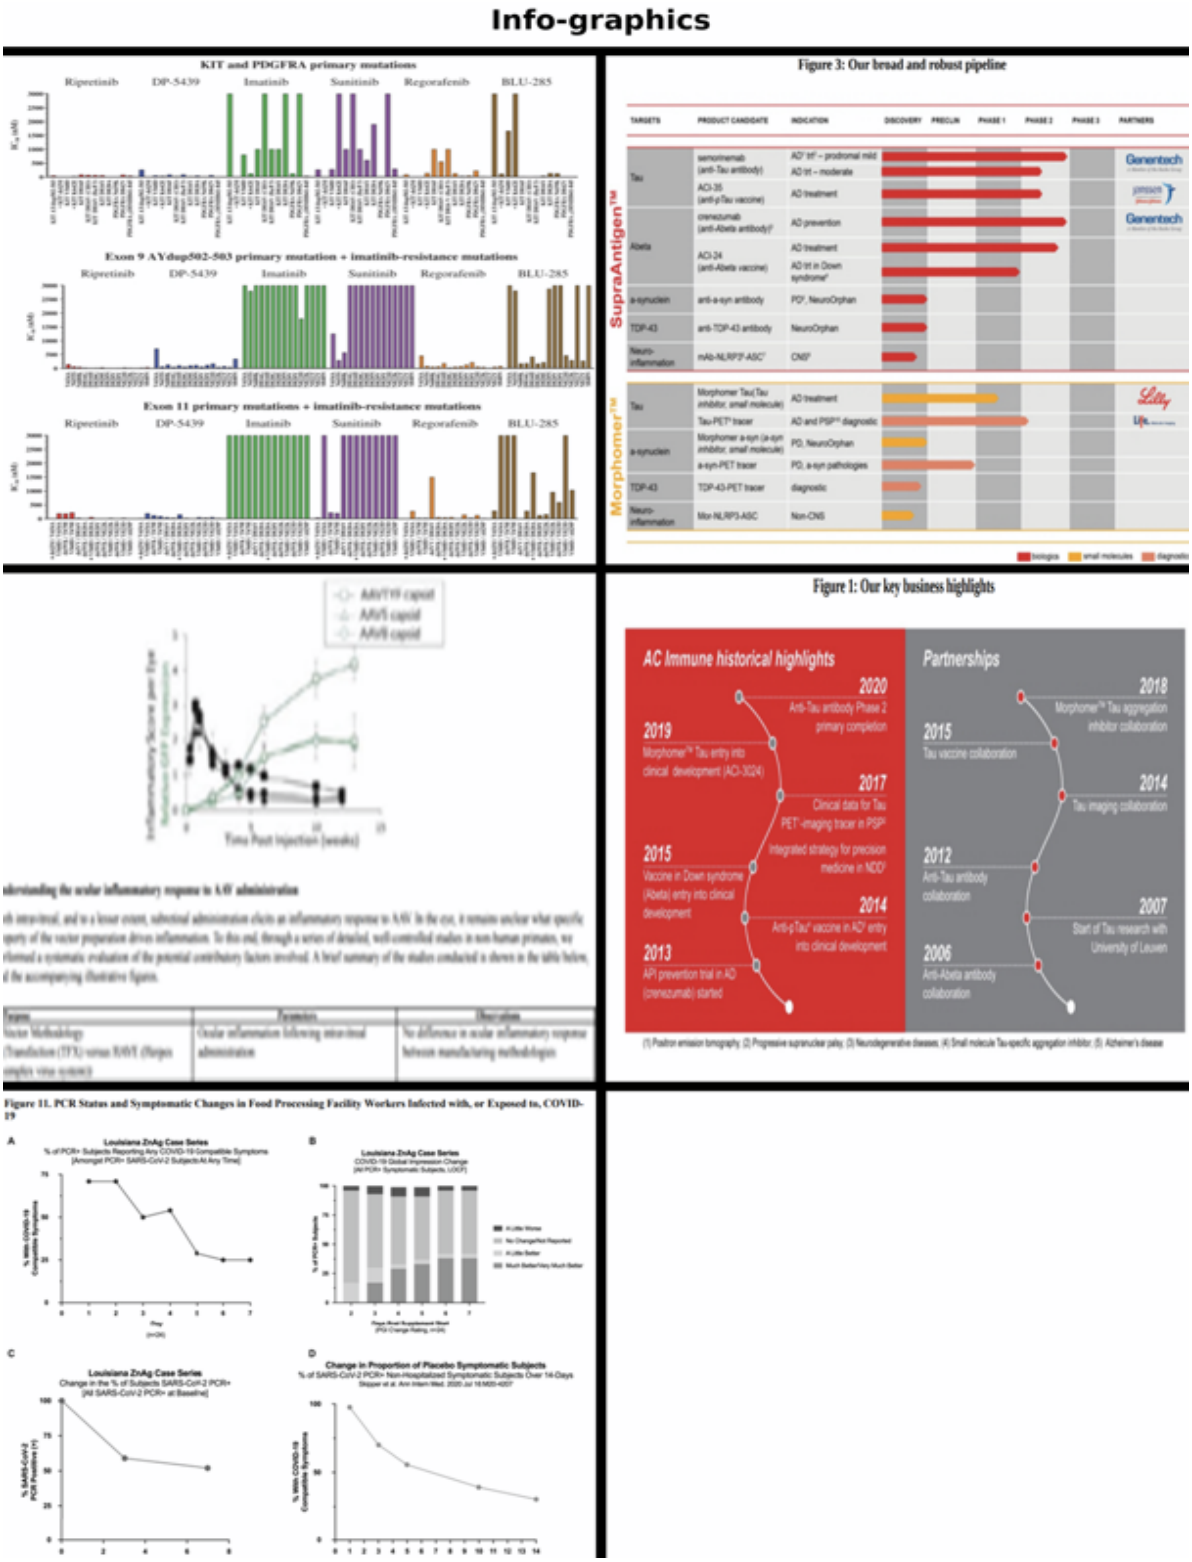
[26] See https://www.sec.gov/divisions/corpfin/guidance/regs-tinterp.htm.

Appendix B. The Categorization of the Keywords from Image Captions

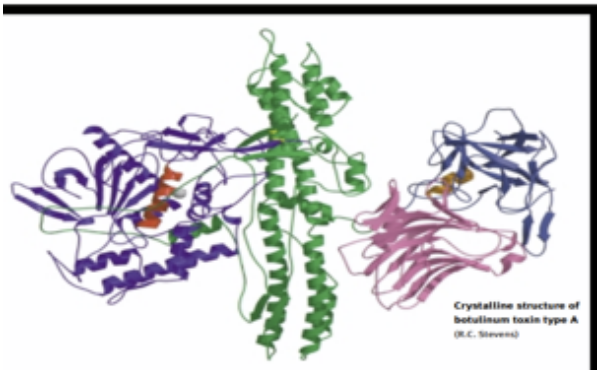
The table lists the categorization of the keywords extracted from image captions in annual reports. This provides a structured framework for analyzing the thematic content of corporate visuals in financial disclosures.

Category	**Keywords**
info-graphics	line graph, diagrams, bar chart, pie chart, histogram, infographic, plot, timeline, bar graph, number of countries, number of states
Logo	logo, brand, poster, symbol
Sceneries	forest, tree, lake, sky, plane, leaves, branches, flowers, field, mountain, snow, skyline, skyscraper, city, building, plant, house, park, cliff, room, wall, street, tunnel, bridge, factory, store, station, restaurant, stadium, swimming pool, tower, shopping center, desert, ocean, sea, river, statue, grass, site, mall, road, aisle, resort, museum, hill, hotel
People	man, woman, baby, babies, child, guy, girl, boy, skateboarder, doctor, engineer, person, father, mother, daughter, son, family, worker
Teams and Groups	men, women, people, crew, group of, team
Map	map
Hi-Tech and Innovation	3d print, 5g, actuator, aerial, aerospace, algorithm, astronaut, automation, composite, computing, conductor, cpu, data storage, development, drone, engine, energy, high speed, high-speed, innovation, infrastructure, intelligence, internet, jet, layers of the earth, mars, motherboard, nano, nasa, network, neural, operating light, processor, renewable, robot, rocket, satellite, Saturn, sensor, server, software, solar, spacecraft, spaceship, technology
Transport	car, bicycle, bike, motorcycle, scooter, train, rail, subway, bus, coach, carriage, cart, truck, ship, tanker, boat, cruise, yacht, ferry, kayak, drone, helicopter, air balloon
Farming and Agricultural	crops, livestock, irrigation, fertilizer, seed, tractor, barn, diary, orchard, pesticide, herbicide, grain, vegetable, compost, weed, manure, tiling, agriculture
Healthcare and Bio-tech	antibiotics, antibody, antigen, biological, blood, blood pressure, cell, cervical spine, chromosome, cotton swab, dna, doctor, drug, gene, genome, gel, gums, hand cream, hospital, immune system, laboratory, lung, mask, medical, molecule, ointment, organ, oxygen, pads, pain relievers, patient, pharmaceutical, prescription, pills, protein, rna, scalpel, scissors, stethoscope, spray, surgical, surgical light, syringe, test tube, teeth, tumor, toothpaste, vaccine, ventilator, vitamin, viral, vision, wipe
Food/Grocery and Lifestyle	tomato, taco, broccoli, carrots, spinach, banana, apple, milk, diary, egg, butter, chicken, beef, fillet, shrimp, bacon, rice, pasta, bread, sandwich, oats, bean, ice cream, berries, fries, chips, coffee, tea, juice, soda, snack, chocolate, soap, wine, beer, alcohol, bottle
Home Item	teapot, vases, carpet, cloth, sofa, table, chair, bed, dresser, curtain, towel, microwave, fridge, washing machine, cleaner, television, mirror, clock, rug, kitchen, cup, camera, ball, dishwasher, refrigerator, phone, iphone, galaxy s, galaxy tab, ipad, cell phone, smartphone, air condition
Abstract Image	abstract

Appendix C. The Example of Categorized Images



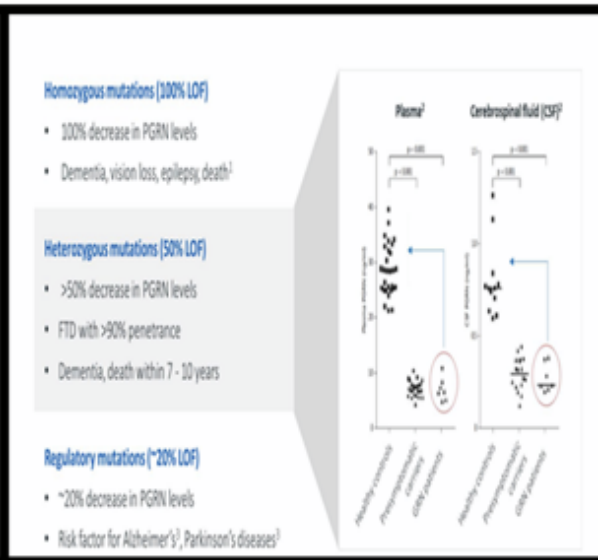
Healthcare and Biotech

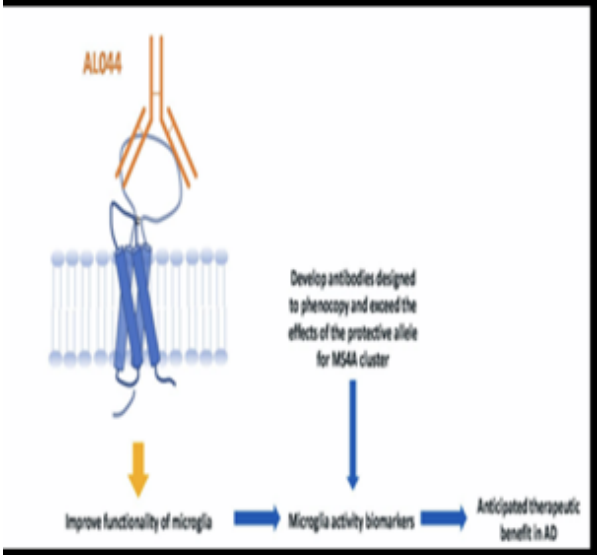


Crystalline structure of botulinum toxin type A (R.C. Stevens)

Aesthetic and Skin Care Market

The interest in aesthetic and skin care uses of BOTOX has increased considerably. In late 2001, the Company expects regulatory approval to market BOTOX for glabellar lines in the United States, Canada, and at least one major European market. Full commercial launch activities are planned, with Allergan leveraging its existing strengths in dermatology and skin care. Allergan's U.S. dermatology sales team,





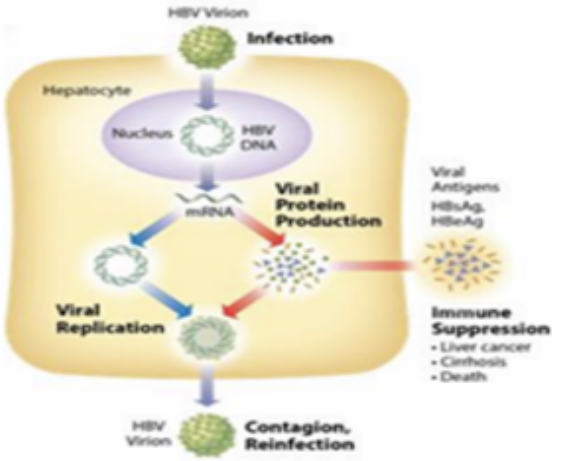
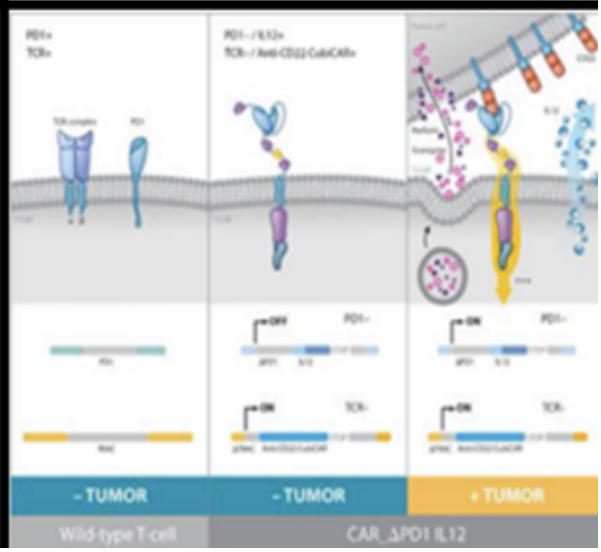
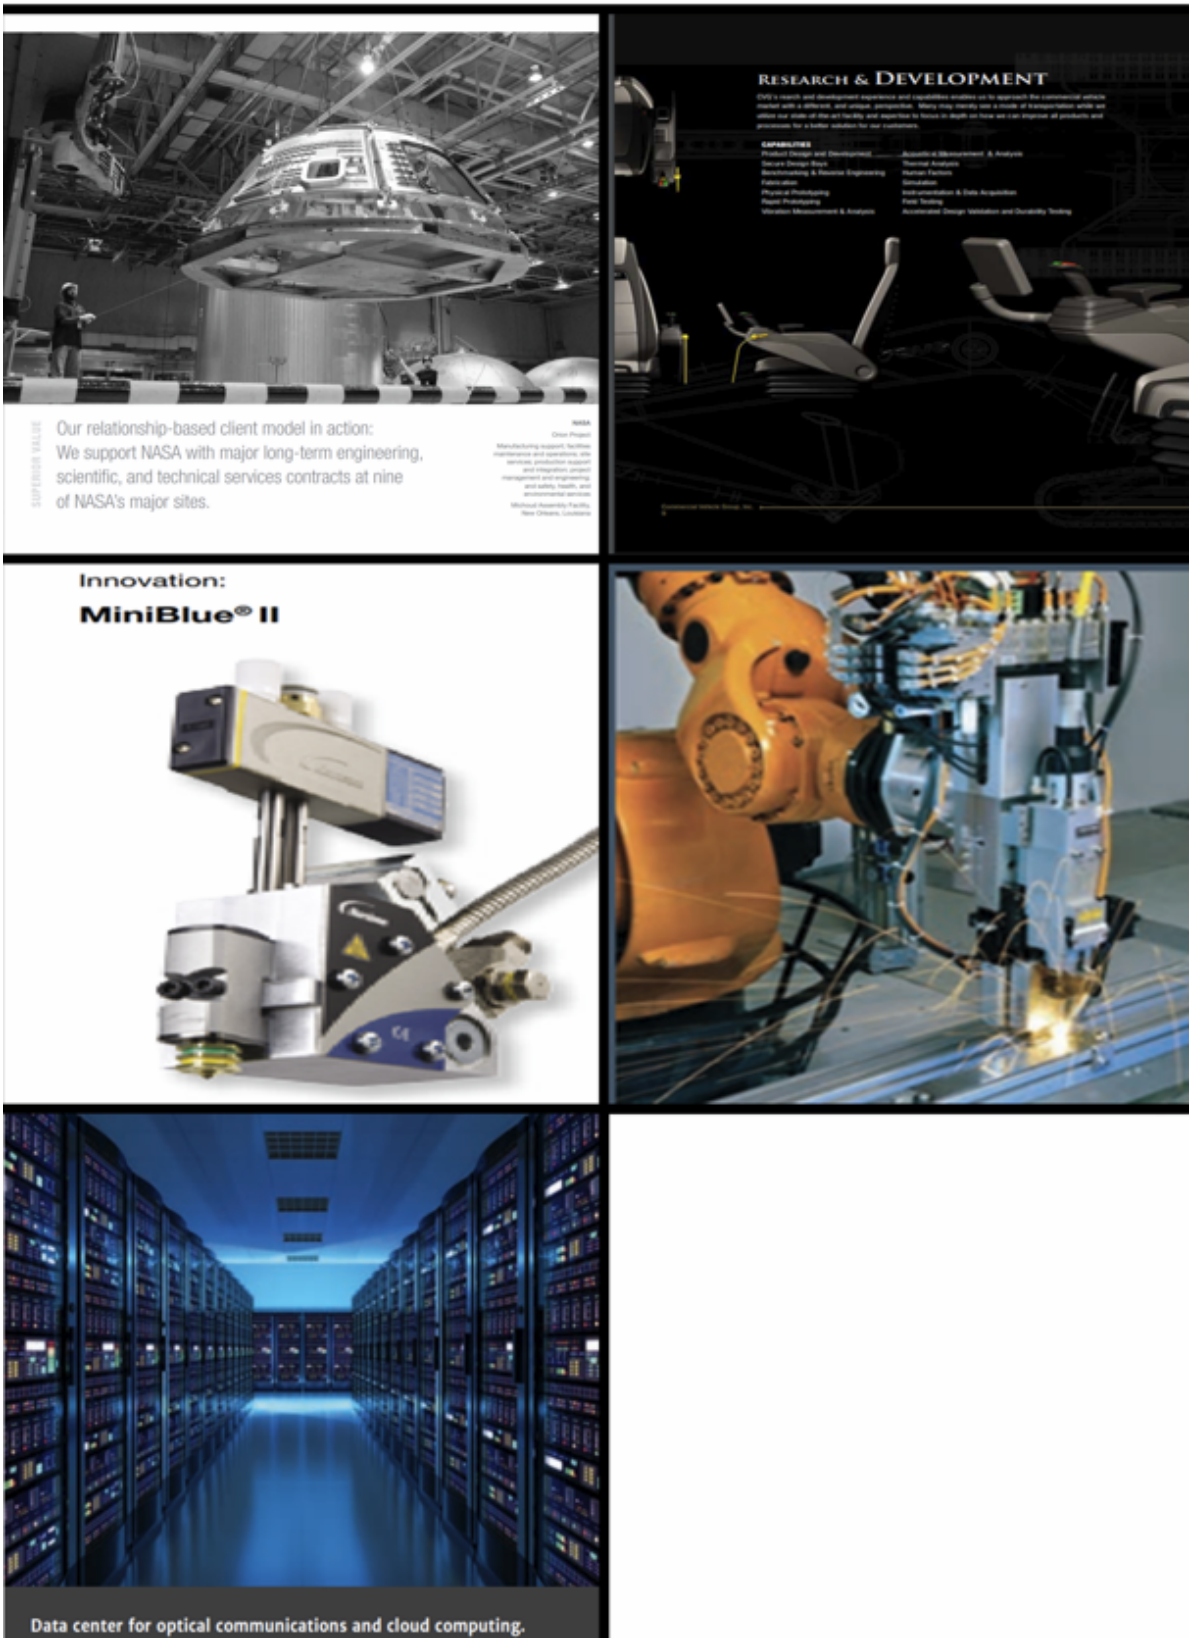


Figure 2: Chronic HBV mechanism untreated

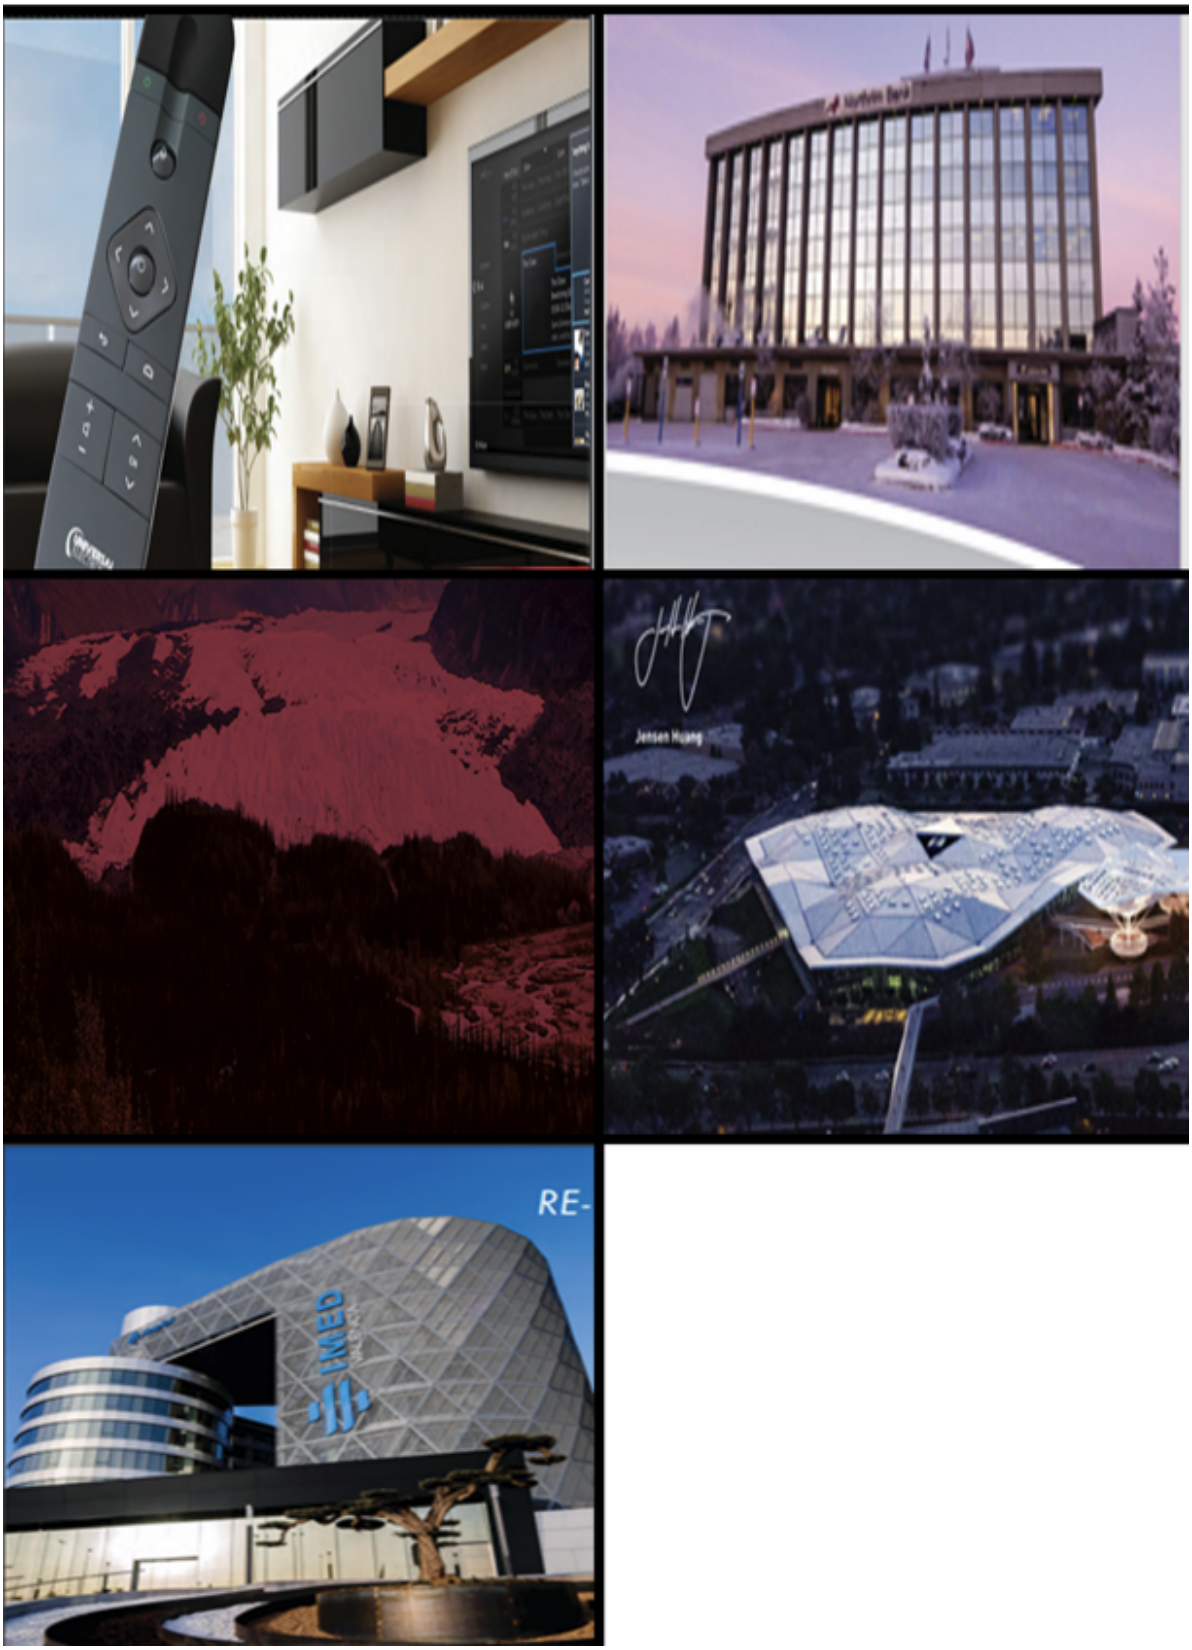




Hi-Tech and Innovation



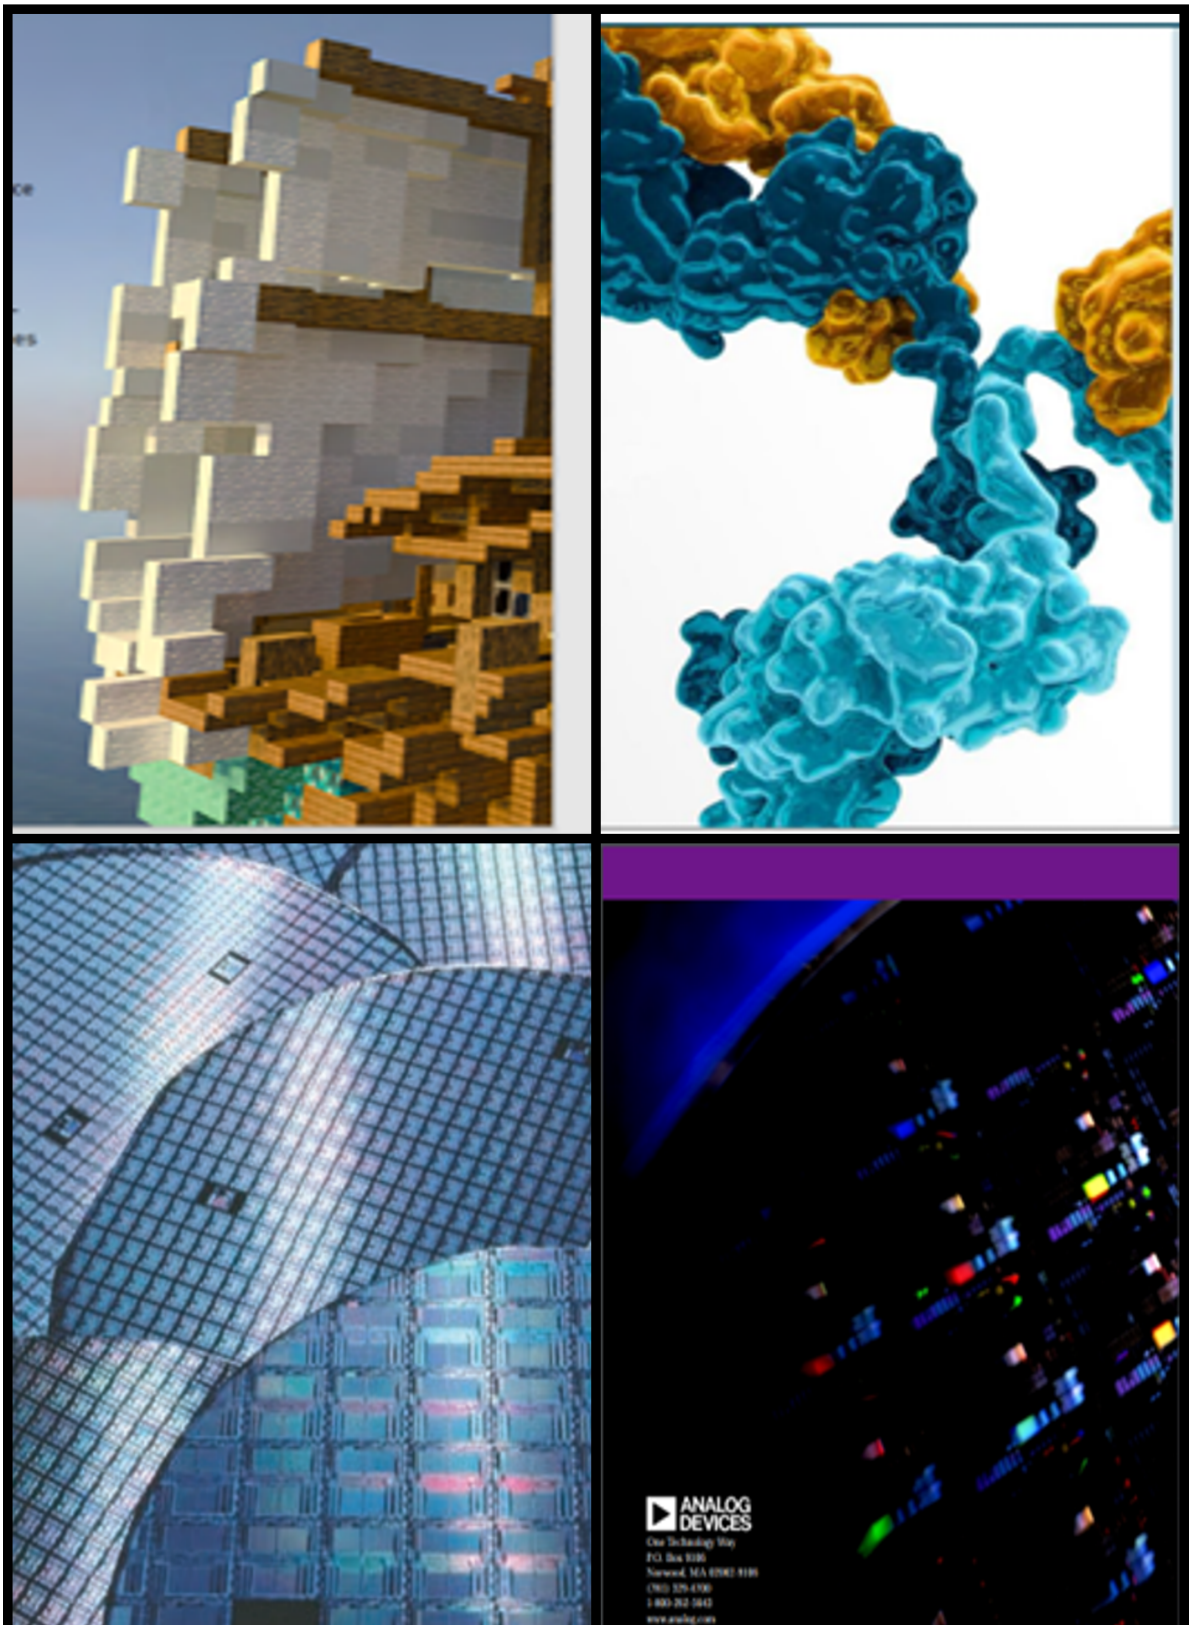
Scenery



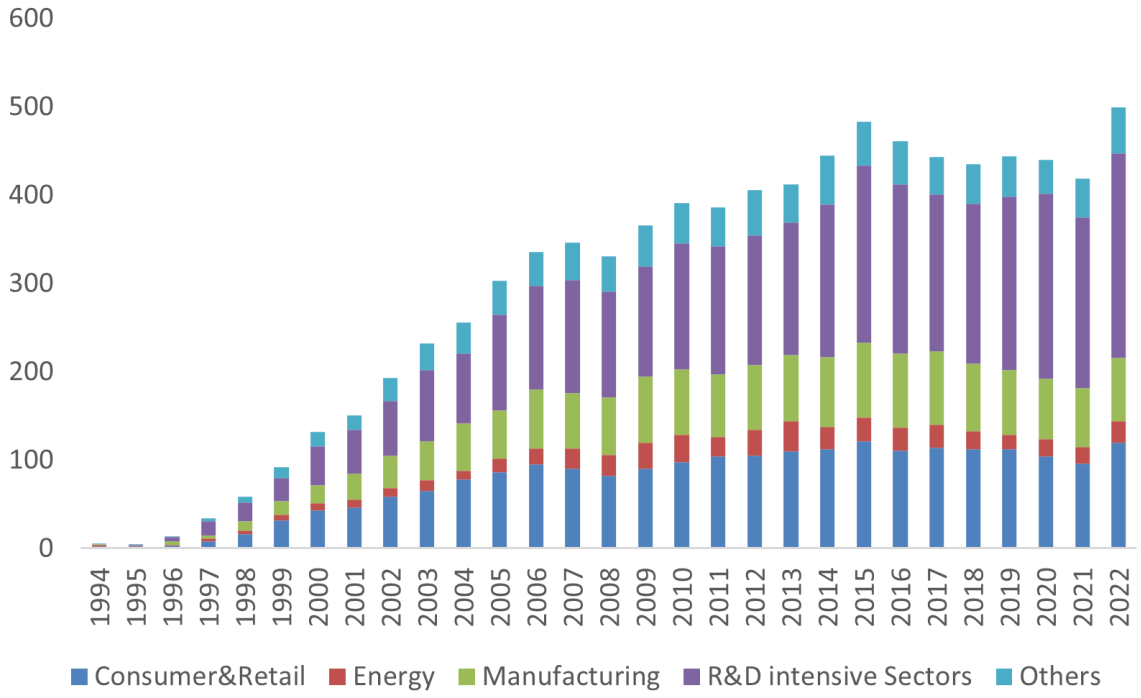
People


PAPER because
learning something new deserves all your attention.

Abstract Images



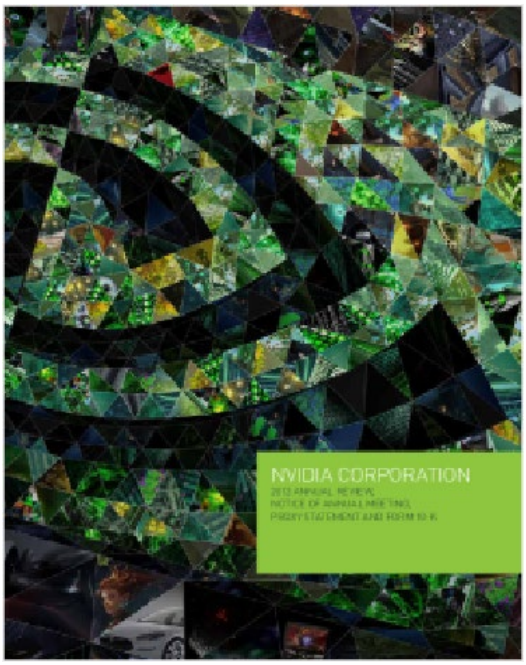
Appendix D. The Use of Visuals Across Different Sectors



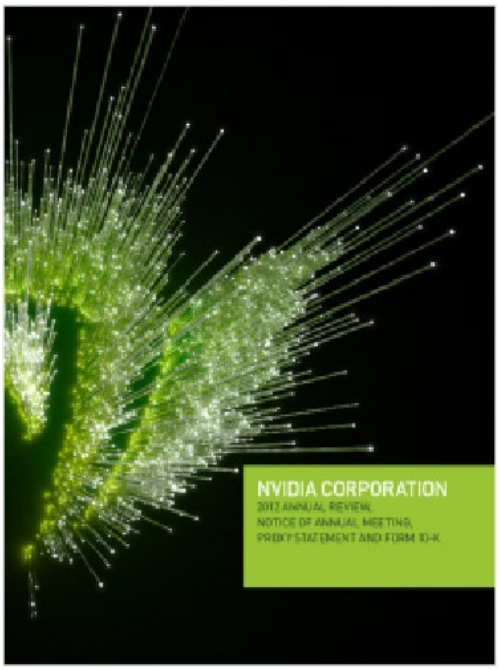
This figure shows the number of firms using visuals in their reports in each year, categorized by sector. The sectors include Consumer&Retail, Energy, Manufacturing, R&D-intensive Sectors, and Other, from 1994 to 2022.

Appendix E. The Example of Nvidia's Annual Reports from 2010 to 2013

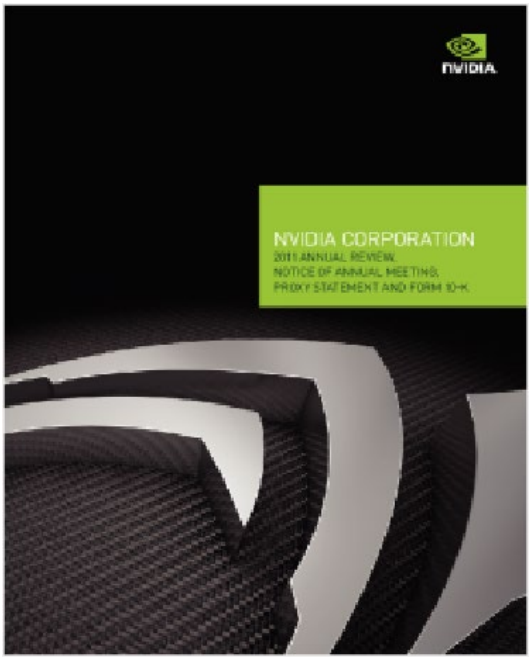
2013



2012



2011



2010



Appendix F. Cover Page of Nvidia's Annual Report/10-K Issued in 2009

UNITED STATES
SECURITIES AND EXCHANGE COMMISSION
Washington, D.C. 20549

FORM 10-K

[x] **ANNUAL REPORT PURSUANT TO SECTION 13 OR 15(d) OF THE SECURITIES EXCHANGE ACT OF 1934**
For the fiscal year ended January 25, 2009

OR

[_] **TRANSITION REPORT PURSUANT TO SECTION 13 OR 15(d) OF THE SECURITIES EXCHANGE ACT OF 1934**

Commission file number: 0-23985

NVIDIA CORPORATION

(Exact name of registrant as specified in its charter)

Delaware	**94-3177549**
(State or other jurisdiction of Incorporation or Organization)	**(I.R.S. Employer Identification No.)**

2701 San Tomas Expressway
Santa Clara, California 95050
(408) 486-2000
(Address, including zip code, and telephone number, including area code, of principal executive offices)

Securities registered pursuant to Section 12(b) of the Act:

Title of each class	**Name of each exchange on which registered**
Common Stock, $0.001 par value per share	The NASDAQ Global Select Market

Securities registered pursuant to Section 12(g) of the Act:
None

Indicate by check mark if the registrant is a well-known seasoned issuer, as defined in Rule 405 of the Securities Act.
Yes ☒ No ☐

Indicate by check mark if the registrant is not required to file reports pursuant to Section 13 or Section 15(d) of the Act.
Yes ☐ No ☒

Indicate by check mark whether the registrant (1) has filed all reports required to be filed by Section 13 or 15(d) of the Securities Exchange Act of 1934 during the preceding 12 months (or for such shorter period that the registrant was required to file such reports), and (2) has been subject to such filing requirements for the past 90 days.
Yes ☒ No ☐

Indicate by check mark if disclosure of delinquent filers pursuant to Item 405 of Regulation S-K (§ 229.405 of this chapter) is not contained herein, and will not be contained, to the best of registrant's knowledge, in definitive proxy or information statements incorporated by reference in Part III of this Form 10-K or any amendment to this Form 10-K. ☒

Indicate by check mark whether the registrant is a large accelerated filer, an accelerated filer, a non-accelerated filer, or a smaller reporting company. See definitions of "large accelerated filer", "accelerated filer" and "smaller reporting company" in Rule 12b-2 of the Exchange Act. (Check one):

Large accelerated filer ☒ Accelerated filer ☐

Non-accelerated filer ☐ (Do not check if a smaller reporting company) Smaller reporting company ☐

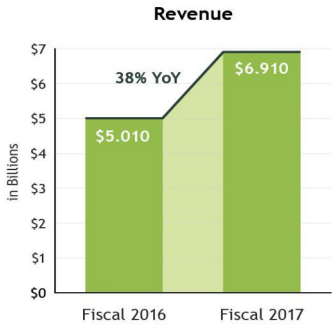
Indicate by check mark whether the registrant is a shell company (as defined in Rule 12b-2 of the Act).

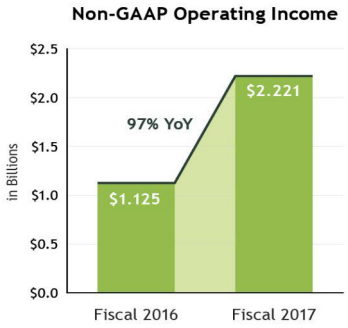
Appendix G. Selected PRO-Images about Nvidia's Products and Services

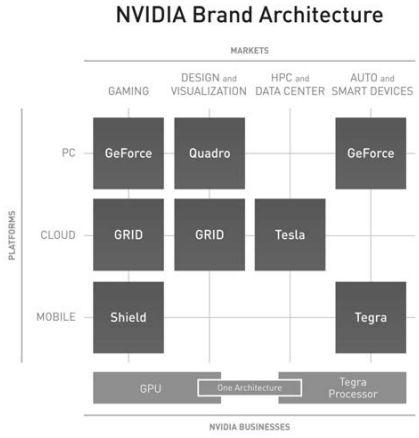


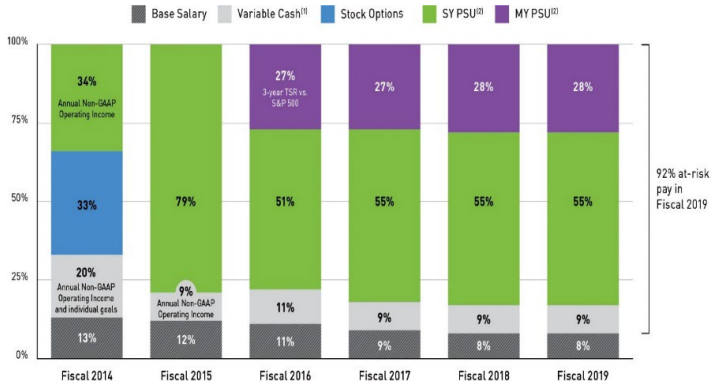
Appendix H. Selected Info-graphics about Performance and Compensations

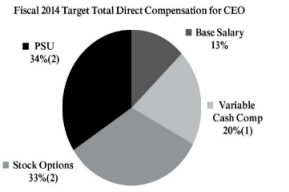


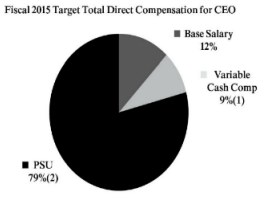


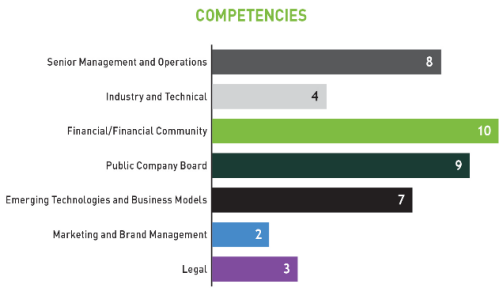


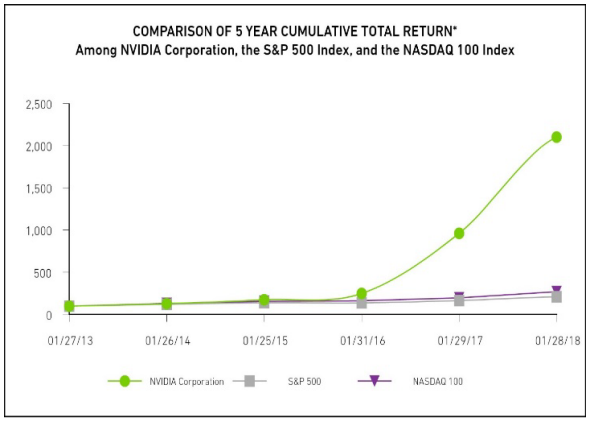


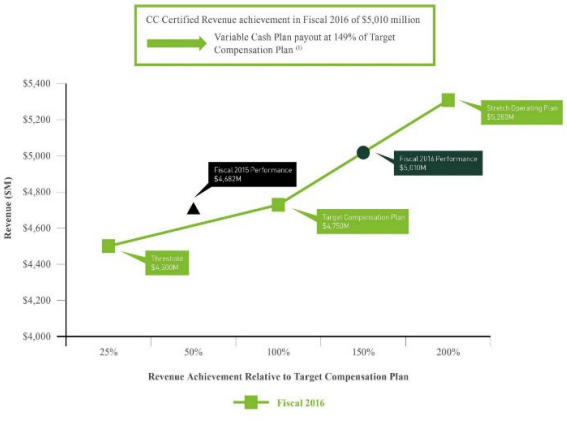












Appendix I. Proof of Proposition 1

First, we consider the ineffective case where $\mu = 0$, i.e., the manager's efforts have no impacts on any attentive traders. In this case, the manager will always choose $s = 0$ to save cost and all market participants will be ignorant of the component $\tilde{\alpha}$ until it is revealed. Given $s = 0$, the trading dynamics in our model is identical to a two-period special case of Holden and Subrahmanyam (1992) who characterize the subgame perfect linear equilibrium for the general multi-period case with multiple traders. Here, we provide analytical solution for the two-period case, following the general procedure of Proposition 1 in Holden and Subrahmanyam (1992). We use the same notation as theirs. Let $\Sigma_t := \text{Var}(\tilde{v}_0|\tilde{p}_t, \tilde{p}_{t-1}, ...\tilde{p}_0)$ represent the posterior variance of $\tilde{v}_0$ conditional on the price history up to time $t \in \{0, 1, 2\}$. In a linear equilibrium, market makers follow the linear pricing strategy (2), with

$$\lambda_1 = \frac{M\beta_1\Sigma_1}{\sigma_u^2}, \qquad \lambda_2 = \frac{M\beta_2\Sigma_2}{\sigma_u^2}, \tag{I-1}$$

where β_t is the trading intensity of each informed trader at time $t = 1, 2$. Let $\eta := \lambda_2/\lambda_1$. By backward induction, it is easy to derive that the optimal trading strategy for trader i is

$$X_{1,i}(\tilde{v}_0) = \beta_1 \cdot (\tilde{v}_0 - p_0) = \frac{\eta(M+1)^2 - 2}{\eta(M+1)^3 - 2M} \cdot \frac{\tilde{v}_0 - p_0}{\lambda_1} \tag{I-2}$$

$$X_{2,i}(\tilde{v}_0, \tilde{p}_1) = \beta_2 \cdot (\tilde{v}_0 - \tilde{p}_1) = \frac{\tilde{v}_0 - \tilde{p}_1}{\lambda_2(M+1)}. \tag{I-3}$$

The Bayesian update of the posterior variance about $\tilde{v}_0$ follows

$$\Sigma_2 = (1 - M\beta_2\lambda_2)\Sigma_1 = \frac{\Sigma_1}{M+1}, \qquad \Sigma_1 = (1 - M\beta_1\lambda_1)\Sigma_0 = \frac{\eta(M+1)^2\sigma_v^2}{\eta(M+1)^3 - 2M}. \tag{I-4}$$

Combining (I-1), (I-2), (I-3), and (I-4), we can solve the expressions for λ_1 and λ_2:

$$\lambda_1 = \frac{\sqrt{\eta M(M+1)^2(\eta(M+1)^2 - 2)}}{\eta(M+1)^3 - 2M} \cdot \frac{\sigma_v}{\sigma_u}, \qquad \lambda_2 = \eta\lambda_1 = \sqrt{\frac{\eta M}{\eta(M+1)^3 - 2M}} \cdot \frac{\sigma_v}{\sigma_u}, \tag{I-5}$$

Define the ratio of lambdas as $\eta := \lambda_2/\lambda_1$. Then Equation (I-5) implies

$$\frac{\lambda_2^2}{\lambda_1^2} = \frac{\eta(M+1)^3 - 2M}{\eta(M+1)^4 - 2(M+1)^2} = \eta^2 \tag{I-6}$$

This identity leads to the cubic equation, $(M+1)^4\eta^3 - 2(M+1)^2\eta^2 - (M+1)^3\eta + 2M = 0$. For $M \geq 1$, there exists a unique solution that meets the second order condition and the constraint $\lambda_t > 0$ for $t = 1, 2$. Stock prices are not efficient because market markers have overlooked a volatile component, $\tilde{\alpha}$, in the firm value. Traders in this case are only *partially* informed: none of them are aware of alpha as it cannot be effectively conveyed ($\mu = 0$). Nonetheless, stock prices are unpredictable based on the public information because all participants in financial markets have the common knowledge that the firm's liquidation value is drawn from the normal distribution $\mathcal{N}(1, \sigma_v^2)$. Therefore, the absence of abnormal returns is not indicative of market efficiency. It can result from ineffective disclosures ($\mu = 0$).

Now we consider the general, nontrivial case where $\mu \in (0, 1]$, i.e., the manager's extra efforts ($s = 1$) are effective to some extent. Note that if the manager chooses to make no efforts ($s = 0$) for whatever reason, traders cannot tell if it is due to no alpha ($\tilde{\alpha} = 0$), overly uncertain alpha ($\sigma_\alpha^2 \gg \overline{\alpha}$), or negative alpha ($\overline{\alpha} < 0$). For fundamental investors (as assumed in our model), they will not engage in speculative trading if they have not found any superior information. Thus, the trading equilibrium when firms choose $s = 0$ will be identical to the one under the ineffective communication ($\mu = 0$) as we have just discussed.

The difference occurs only when the manager makes extra efforts ($s = 1$). Even in this case, the rest of the market is unaware of $\tilde{\alpha}$ until the value of $\tilde{v} = \tilde{v}_0 + \tilde{\alpha}$ is publicly revealed. Therefore, market makers follow the same linear pricing strategy described by Equation (2). Nonetheless, after observing $s = 1$, each attentive trader may discover the value of α with probability μ and fail to find it with probability $1-\mu$. The potential impact of the manager's action on each attentive trader can be concisely written as $s\tilde{\alpha}\mathbf{1}_\mu$ which, conditional on the manager's efforts ($s = 1$), is exactly equal to $\tilde{\alpha}$ with probability μ and equal to zero with probability $1 - \mu$. Based on the assumption of fundamental investing, each attentive trader initiates an independent research after observing $s = 1$, which takes at least one period of time before he can find alpha. As fundamental traders, they only respond to the fundamental information they possess. Thus, each attentive trader's strategy at $t = 1$ is driven by his initial private information about $\tilde{v}_0$, independent of the manager's action $s = 1$ (as it is not directly informational). This implies that, in the subgame perfect linear equilibrium, the first-period optimal strategy $X_{i,1}$ is still given by Equation (I-2). However, the trading strategy $X_{i,2}$ at $t = 2$ may be affected by the manager's choice of $s = 1$ since traders may find alpha before $t = 2$. This indicates an extra informational term in their second-period trading strategy $X_{i,2}$. This extra term should be linearly proportional to the product $s\tilde{\alpha}\mathbf{1}_\mu$.

For any trader who has found alpha after one period, he remains uncertain about how many other traders have discovered the same information of alpha. Due to the indepen-

dence of $\tilde{v}_0$ and $\tilde{\alpha}$, the optimal strategy can be separated into two parts and written as $X_{i,2}(\tilde{v}_0, \tilde{p}_1, \tilde{s}) = X^o_{i,2}(\tilde{v}_0, \tilde{p}_1) + Z_i(\tilde{s}\tilde{\alpha}\mathbf{1}_\mu)$. The first term is still given by Equation (I-3): $X^o_{i,2}(\tilde{v}_0, \tilde{p}_1) = \frac{\tilde{v}_0 - \tilde{p}_1}{\lambda_2(M+1)}$. The second term, $Z_i(\tilde{s}\tilde{\alpha}\mathbf{1}_\mu)$, is random as it depends on the discovery of alpha. Note that $\tilde{\alpha}$ is unexpected by market markers and the initial price change, $\tilde{p}_1 - p_0$, only reflects the market expectation of $\tilde{v}_0$ conditional on the total order flow $\tilde{y}_1$. Thus, before trading occurs at time $t = 2$, no information about alpha has been incorporated into the prices. If trader i has found the value of alpha after research, then we can write $Z_i(\tilde{s}\tilde{\alpha}\mathbf{1}_\mu) = z_i(\alpha)$ and the alpha-related part of trader i's objective function can be written as

$$\max_{z_i} \mathrm{E}\left[\left(\alpha - \lambda_2\left(z_i + \sum_{j=1}^{\tilde{M}_\alpha} z_j(\alpha) + \tilde{u}_2\right)\right) z_i\right], \tag{I-7}$$

where $\tilde{M}_\alpha$ represents the (random) number of other traders who have discovered the value of alpha after observing the manager's action $s = 1$. This is a binomial random variable, $\tilde{M}_\alpha \sim \mathcal{B}(M-1, \mu)$. From the perspective of trader i who has found alpha, there are at most $M - 1$ independent competitors and each of them has the same success rate μ. Equation (I-7) is essentially a one-period problem for a random number of informed traders who compete on the same alpha. Since we are considering a symmetric equilibrium where $z_i(\alpha) = z_j(\alpha) = z(\alpha)$, the first order condition of (I-7) leads to

$$\alpha - \lambda_2\left(2z_i + \mathrm{E}\left[\sum_{j=1}^{\tilde{M}_\alpha} z_j(\alpha)\right]\right) = \alpha - \lambda_2(2z(\alpha) + \mu(M-1)z(\alpha)) = 0, \tag{I-8}$$

which leads to the solution

$$z(\alpha) = \frac{\alpha}{\lambda_2(2 + \mu(M-1))}. \tag{I-9}$$

For $M = 1$, we obtain the simple monopolistic result $z(\alpha) = \frac{\alpha}{2\lambda_2}$ for the one-period problem. For $M > 1$, the uncertainty of the number of informed traders can reduce their competition. In summary, the optimal trading strategy for each attentive trader at $t = 2$ is given by

$$\begin{aligned} X_{i,2}(\tilde{v}_0, \tilde{p}_1, \tilde{s}) &= \frac{\tilde{v}_0 - \tilde{p}_1}{\lambda_2(M+1)} + \tilde{s}\mathbf{1}_\mu z(\tilde{\alpha}) \\ &= \frac{\tilde{v}_0 - \tilde{p}_1}{\lambda_2(M+1)} + \frac{\tilde{s}\tilde{\alpha}\mathbf{1}_\mu}{\lambda_2(2 + \mu(M-1))}, \end{aligned} \tag{I-10}$$

where the first term is driven by the information of $\tilde{v}_0$ and the second term is driven by $\tilde{s}\tilde{\alpha}\mathbf{1}_\mu$. Again, both terms are based on material private information about firm fundamentals. There is no speculative trading on the manager's choice of s at either $t = 1$ or $t = 2$.

Appendix J. Proof of Proposition 2

With rational expectations about the trading equilibrium in Proposition 1, the manager needs to evaluate the benefit from her choice of s against the associated cost. Although the manager anticipates how each trader may respond to her choice of $s \in \{0, 1\}$, she does not know exactly how many traders would find alpha following her action $s = 1$. Based on the result of subgame perfect linear equilibrium and the property that $\mathrm{E}[\tilde{v}_0 - p_0] = \mathrm{E}[\tilde{v}_0 - \tilde{p}_1] = 0$ (because the prices are efficient for $\tilde{v}_0$), the manager's objective (1) becomes

$$\begin{aligned} & \max_{s\in\{0,1\}} \mathrm{E}^{\mathcal{F}}[\tilde{p}_2 - p_0] - s \cdot c = \max_{s\in\{0,1\}} \mathrm{E}^{\mathcal{F}}[\lambda_1 \tilde{y}_1 + \lambda_2 \tilde{y}_2] - s \cdot c \\ = & \max_{s\in\{0,1\}} \mathrm{E}^{\mathcal{F}}\left[\lambda_2 \sum_{i=1}^{M} \frac{s\tilde{\alpha}\mathbf{1}_\mu}{\lambda_2(2+\mu(M-1))}\right] - s \cdot c \\ = & \max_{s\in\{0,1\}} \frac{s\mu M\overline{\alpha}}{2+\mu(M-1)} - s \cdot c \end{aligned} \tag{J-11}$$

This leads to her binary decision rule with the threshold, $K = \frac{2+\mu(M-1)}{\mu M}c$. When $\overline{\alpha} \geq K$, the manager chooses $s = 1$ and the stock will experience no abnormal return over the time window $[0, 1]$ but positive abnormal return $\mathrm{E}[\tilde{p}_2 - p_0] = \frac{\mu M\overline{\alpha}}{2+\mu(M-1)}$ afterwards. When $\overline{\alpha} < K$, the manager chooses $s = 0$ such that there will be no abnormal returns up to either $t = 1$ or $t = 2$, that is, $\mathrm{E}[\tilde{p}_1 - p_0 \mid s = 0] = \mathrm{E}[\tilde{p}_2 - p_0 \mid s = 0] = 0$, follows from the property that, in the subgame perfect linear equilibrium, market markers set efficient prices for $\tilde{v}_0$, i.e., $\tilde{p}_1 = \mathrm{E}[\tilde{v}_0|\tilde{y}_1]$ and $\tilde{p}_2 = E[\tilde{v}_0|\tilde{y}_1, \tilde{y}_2]$. Nonetheless, those prices are not efficient in any case because the market is ignorant of the volatile component $\tilde{\alpha}$.

As a note, we have assumed for simplicity that the manager's objective function is described by Equation (1). This is sufficient for the testable implications. There are different ways to describe the manager's utility. For example, a more complex specification can be

$$\max_{s\in\{0,1\}} \mathrm{E}^{\mathcal{F}}[\delta \cdot (\tilde{p}_2(s) - p_0)] - s \cdot c - \gamma \mathrm{Var}^{\mathcal{F}}[\delta \cdot (\tilde{p}_2(s) - p_0)]. \tag{J-12}$$

where the parameter δ reflects the strength of the manager's incentive to boost her firm's stock price and the parameter γ represents the manager's risk-aversion coefficient. This specification can make the expression of K complicated but not alter our main implications.

Our simple model can be extended or modified in many other ways. For example, the firm manager may be concerned about the time "integral" of price changes; the manager's choice variable s can be continuous over the interval $[0, 1]$. Those model variants do not change the key implications derived from the simple model in the main text.

Appendix K. Definitions of Variables

Variables	**Definition**
Return	Cumulative weekly stock returns for a given event window
Visual Report	Dummy variable equal to 1 if the firms use graphic annual reports in the fiscal year, and 0 otherwise
S_{Images}	Dummy variable equal to 1 if a firm adds a visually enhanced annual report on top of a plain 10-K in a fiscal year, and 0 otherwise
$S_{Info-graphics}$	Dummy variable equal to 1 if a firm adds a visually enhanced annual report, in which includes info-graphics such as charts, tables, plots, and figures, on top of a plain 10-K in a fiscal year, and 0 otherwise
S_{IP}	Dummy variable equal to 1 if a firm adds a visually enhanced annual report, in which includes illustrative and photographic images, on top of a plain 10-K in a fiscal year, and 0 otherwise
S_{ITHB}	Dummy variable equal to 1 if a firm adds a visually enhanced annual report, in which includes info-graphics, in which includes visuals related to advancements in innovation and technology related to IT, healthcare, and biotechnology sectors (i.e., R&D-focused visuals), on top of a plain 10-K in a fiscal year, and 0 otherwise
Readability	10-K FOG readability index
Attention	The aggregated number of Bloomberg institutional attention shocks within a given time window
INST	Institutional ownership at the beginning of a fiscal year
Institutional Flow	Institutional net inflow in percentage, by taking the difference between the newly available institutional ownership and the reference institutional ownership
ROE	Return on equity at the beginning of a fiscal year
Leverage	Long term debt (DLTT) plus debt in current liabilities (DLC) scaled by the sum of long term debt, debt in current liabilities, and total stockholders' equity (SEQ) $\times$ 100, at the beginning of a fiscal year
Size	Natural log of total assets at the beginning of a fiscal year
R&D	Research and Development Expenses (XRD) scaled by end-of-year total assets (AT) $\times$ 100
Firm Age	Natural log of firm age
Patents	The number of patents filed that are eventually granted during a fiscal year
New Products	The number of new product introductions during a fiscal year

Appendix L. Supplementary Analysis

Table L1: Firm and industry Characteristics and the Choice of Using Visuals

This table reports the results from estimating the following logistic regression, which identify the characteristics of firms that use graphic elements in their annual reports:

$$\begin{aligned} Logit(Visuals_{i,t}) \;=\; & \beta_0 + \beta_1 \times Size_{i,t-1} + \beta_2 \times ROE_{i,t-1} + \beta_3 \times Tobin's\ Q_{i,t-1} \\ & + \beta_4 \times 3Y\ Readability_{i,t-1} + \beta_5 \times Firm\ Age_{i,t} \\ & + \beta_6 \times R\&D\ Sectors_{i,t} + \epsilon_{i,t}. \end{aligned}$$

Columns (1) through (6) report the results of Size, ROE, Tobin's Q, Firm Age, and past 3-year average Readability Index, and the indicator of R&D sectors as the independent variable without controlling for other firm and industry characteristics. Column (7) reports the results by including all the above variables as the independent variables. We report the t-statistic in the parenthesis under each coefficient. *, **, *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	S_{Images}	S_{Images}	S_{Images}	S_{Images}	S_{Images}	S_{Images}	S_{Images}
Size	0.037*						0.001
	(1.87)						(0.04)
ROE		-0.001					-0.001
		(-0.93)					(-0.36)
Tobin's Q			0.011				0.002
			(0.69)				(0.10)
3Y Readability				-0.028**			-0.028**
				(-2.41)			(-2.34)
Firm Age					-0.032		-0.121**
					(-0.97)		(-2.12)
R&D Sectors						0.164**	0.160*
						(2.46)	(1.73)
N	11,978	10,150	9,409	7,295	10,104	11,978	6,758
pseudo R-sq	0.001	0.000	0.000	0.001	0.000	0.001	0.004