Dissecting Corporate Culture Using Generative AI – Insights from Analyst Reports*

Kai Li, Feng Mai, Rui Shen, Chelsea Yang, and Tengfei Zhang

First version: August 2023 This version: February 2024

Abstract

Our study is among the first in finance, accounting, and economics to apply generative AI models as reasoning agents on analyst reports to gain insights into analysts' views of corporate culture. We employ generative AI (ChatGPT) to analyze 2.4 million analyst reports between 2000 and 2020. Generative AI organizes analysts' views into a knowledge graph that links different cultural values to their perceived causes and effects. In terms of influencing factors for corporate culture, analysts identify business strategy as a key factor for a large number of cultural values – customer-oriented, operations-oriented, innovation, results-oriented, adaptability, integrity, and risk control; and management team for teamwork, innovation, results-oriented, adaptability, integrity, risk control, and people-oriented values. In terms of business outcomes shaped by corporate culture, analysts identify innovation and adaptability as affecting almost all aspects of business operations, ranging from market share and growth to corporate ESG practices, while other values such as customer-oriented, operation-oriented, or people-oriented have less impact on business outcomes. We further provide evidence that analysts' views of culture are distinct from values presented on corporate websites, and/or from the views of executives and employees. Finally, we show that analysts' views of corporate culture are reflected in their stock recommendations and target price forecasts as well as impact price reactions to the release of their reports. We conclude that analysts' research on corporate culture offers new insights into its causes and effects, and that we are closer to establishing the culture-firm value link.

Keywords: corporate culture; equity analysts; analyst reports; large language models; generative AI; ChatGPT; cause-effect knowledge graph

JEL classifications: C45; C55; G32; G34; Z1

---

* Kai Li: Sauder School of Business, University of British Columbia. Email: kai.li@sauder.ubc.ca. Feng Mai: School of Business, Stevens Institute of Technology. Email: feng.mai@stevens.edu. Rui Shen: Shenzhen Finance Institute, School of Management and Economics, The Chinese University of Hong Kong, Shenzhen. Email: ruishen@cuhk.edu.cn. Chelsea Yang: Sauder School of Business, University of British Columbia. Email: chunxue.yang@sauder.ubc.ca. Tengfei Zhang: Rutgers School of Business-Camden, Rutgers University. Email: tengfei.zhang@rutgers.edu. We thank helpful comments from Hoa Briscoe-Tran, Jin-Chuan Duan, Roby Lehavy (discussant), Johan Sulaeman, Baozhong Yang, Bernard Yeung, Frank Yu, seminar participants at CEIBS, Chinese University of Hong Kong, Chinese University of Hong Kong - Shenzhen, Fanhai International School of Finance, Fudan University, Nanyang Technological University, National University of Singapore, NYU – Shanghai, NYU Stern Corporate Governance Research Luncheon, Renmin University, Rutgers University, Shanghai Jiaotong University, Shanghai Advanced Institute of Finance, Singapore Management University, Southern University of Science and Technology, Stevens Institute of Technology, Tsinghua University, and conference participants at the First New Finance Conference (Toronto), the New Idea Session at the Financial Management Association Meetings (Chicago), the HKUST Accounting Research Symposium (Hong Kong), and the 2024 Culture Connect Conference (Berkeley). We acknowledge financial support from the Social Sciences and Humanities Research Council of Canada (Grant Number: 435-2022-0285). Li acknowledges financial support from the Canada Research Chair in Corporate Governance. Zhang acknowledges data support from the Whitcomb Center for Research in Financial Services at Rutgers University. All errors are our own.

Dissecting Corporate Culture Using Generative AI – Insights from Analyst Reports

Abstract

Our study is among the first in finance, accounting, and economics to apply generative AI models as reasoning agents on analyst reports to gain insights into analysts' views of corporate culture. We employ generative AI (ChatGPT) to analyze 2.4 million analyst reports between 2000 and 2020. Generative AI organizes analysts' views into a knowledge graph that links different cultural values to their perceived causes and effects. In terms of influencing factors for corporate culture, analysts identify business strategy as a key factor for a large number of cultural values – customer-oriented, operations-oriented, innovation, results-oriented, adaptability, integrity, and risk control; and management team for teamwork, innovation, results-oriented, adaptability, integrity, risk control, and people-oriented values. In terms of business outcomes shaped by corporate culture, analysts identify innovation and adaptability as affecting almost all aspects of business operations, ranging from market share and growth to corporate ESG practices, while other values such as customer-oriented, operation-oriented, or people-oriented have less impact on business outcomes. We further provide evidence that analysts' views of culture are distinct from values presented on corporate websites, and/or from the views of executives and employees. Finally, we show that analysts' views of corporate culture are reflected in their stock recommendations and target price forecasts as well as impact price reactions to the release of their reports. We conclude that analysts' research on corporate culture offers new insights into its causes and effects, and that we are closer to establishing the culture-firm value link.

Keywords: corporate culture; equity analysts; analyst reports; large language models; generative AI; ChatGPT; cause-effect knowledge graph

JEL classifications: C45; C55; G32; G34; Z1

1. Introduction

Since "the 'cultural revolution' in finance" began a decade ago (Zingales 2015), researchers have gained a deeper understanding of corporate culture gleaned from values presented on corporate websites, and through surveys and interviews with corporate executives, earnings conference calls, employee reviews, and job postings (Guiso, Sapienza, and Zingales 2015; Grennan 2019; Li, Mai, Shen, and Yan 2021; Graham, Grennan, Harvey, and Rajgopal 2022a, 2022b; Pacelli, Shi, and Zou 2022; Li, Chen, and Shen 2024); and using proxies for corporate culture (see, for example, Ahern, Daminelli, and Fracassi 2015; Liu 2016; also see surveys by Gorton, Grennan, and Zentefis 2022; Grennan and Li 2023). This body of work as a whole, however, leaves a number of important research questions unanswered; for example, do capital market participants (e.g., information intermediaries) share the same views of corporate culture as corporate insiders (e.g., management)? Does corporate culture affect stock prices? Our study is among the first in finance, accounting, and economics to apply generative artificial intelligence (AI) models as reasoning agents on analyst reports to answer these questions.

Sell-side equity analysts play an important role in processing, producing, and disseminating information about publicly listed companies to capital market participants (see, for example, Womack 1996; Brav and Lehavy 2003; Asquith, Mikhail, and Au 2005; Derrien and Kecskés 2013; Huang, Lehavy, Zang, and Zheng 2018; Birru, Gokkaya, Liu, and Stulz 2022). These analysts gain in-depth knowledge of the firms they follow through both formal and informal channels: reading financial statements, attending earnings conference calls, conducting site visits, and engaging directly with top and divisional managers (Soltes 2014; Brown, Call, Clement, and Sharp 2015). A number of prior papers show that analysts possess value-relevant non-financial information about the companies they cover, for example, management quality and innovation (Previts, Bricker, Robinson, and Young 1994; Huang, Zang, and Zheng 2014; Brown et al. 2015; Bellstam, Bhagat, and Cookson 2021). It is thus natural for us to investigate whether and how analysts conduct analyses on corporate culture, given their information intermediary role. Our research addresses the following questions from the vantage points of 2.4 million analyst reports covering S&P 1500 firms over a 21-year period: 1) What cultural values prevail in a modern corporation? 2) What events, people, and/or systems shape corporate culture? 3) How does a cultural value affect different business outcomes? 4) What is the relationship between corporate culture and price formation?

An analyst report contains quantitative analyses of recent and estimated firm performance (e.g., earnings per share) and qualitative interpretations of information signals (e.g., management quality) (Asquith, Mikhail, and Au 2005; Soltes 2014). An average report comprises approximately 60 sentences over eight pages (Huang, Zang, and Zheng 2014). With millions of reports over our sample period 2000–2020, manually coding reports in search of answers to our research questions is infeasible.

To automatically synthesize analysts' views of corporate culture and delineate the underlying logic in their respective analyses, we introduce a novel method built on generative AI models (e.g., OpenAI's ChatGPT and Google's PaLM). Specifically, we view those models as reasoning agents, capable of extracting causes and effects pertaining to a cultural value from analyst reports (Blanco, Castell, and Moldovan 2008; Radinsky, Davidovich, and Markovitch 2012; Heindorf et al. 2020). These relations reflect the reasoning processes analysts undertake to dissect corporate culture, connecting different cultural values (e.g., innovation and teamwork) to their perceived causes and effects (e.g., management turnover and customer satisfaction). We then use the output from our method to answer the research questions listed above.

We apply generative AI to over 92,000 culture-related segments identified from our report sample. The following is an illustrative example generated from our method. In the report featuring Cisco Systems Inc. by Walter Piecyk from the brokerage firm Painewebber Inc., released on April 14, 2000, the analyst says, *"…Cisco promotes a highly entrepreneurial culture within its organization, which has enabled the company to lure and hire about 3,000 new employees per quarter to its 26,100-employee base.*" Our method extracts a cause-effect relation in a triple as *('innovation culture,' 'enables,' 'ability to lure and hire new employees'*). After converting extracted causes, cultural values, and effects into a normalized format – a process known as canonicalization – we end up with nine cultural values, seventeen drivers of cultural changes, and sixteen business outcomes shaped by corporate culture.

We are aware that extracting analysts' views of corporate culture using generative AI presents several challenges. One such challenge is the untested nature of AI models in accurately extracting cause-effect relations. Concerns also arise regarding the possibility that analysts might not possess meaningful insights into intangibles such as corporate culture, and/or might simply echo information heard from earnings conference calls. To address these issues, we adopt a multi-pronged approach to validate our method.

First, we conduct a direct validation of generative AI models' extraction results through human annotation. We find that ChatGPT and PaLM models both exhibit high levels of accuracy, precision, recall, and F1. In particular, ChatGPT achieves accuracy rates of 91.0%, 86.5%, 88.5%, 71.0%, and 91.5% for cultural values, causes, effects, cause-effect relations, and tones, respectively, and outperforms PaLM. Given the task's complexity, these performance metrics are indicative of generative AI models' reasoning capabilities. Second, we show that analysts' discussions of certain values such as customer-oriented, results-oriented, and risk control cultures are more intense in firm-year observations that do not hold calls compared to those in firm-year observations that hold calls, suggesting that analysts are not simply repeating what they hear from calls. Moreover, we note that even in firm-year observations in which executives say little about corporate culture in calls, analysts covering those firm-years still discuss more about risk control culture than their peers covering firm-year observations in which executives say more about corporate culture in calls. We further note that the nine cultural values featured in analyst reports do not directly repeat those values listed on corporate websites (Guiso, Sapienza, and Zingales 2015), discussed in calls (Li et al. 2021), discussed by employees (Grennan 2019), or discussed by executives in surveys and interviews (Graham et al. 2022a, 2022b), which again suggests that analysts do possess unique insights into corporate culture.

To provide a bird's-eye view of corporate culture from the vantage points of equity analysts, we aggregate the nine cultural values, seventeen different causes, and sixteen different business outcomes into a cause-effect knowledge graph. The graph shows that, according to equity analysts, the top three cultural values prevailing in a modern corporation are innovation, adaptability, and customer-oriented cultures; the top three drivers of cultural changes are business strategy, management team, and strategic transformation; the top three business outcomes that culture shapes are market share and growth, profitability, and employee satisfaction. Our method further allows us to pinpoint the events, people, and/or systems that shape specific cultural values, and identify which cultural values have the most impact in driving specific business outcomes. For example, analysts identify that mergers and acquisitions (M&As, an event) and management team (people) are the top two factors shaping the cultural value of teamwork, and that the cultural value of teamwork bolsters business relationships and drives market share and growth.

We also ask generative AI to classify tones in culture-related segments (as negative, neutral, or positive) and summarize reasons why analysts discuss culture. We find that innovation, customer satisfaction, and employee satisfaction are viewed positively, whereas misconduct, internal conflicts, and risk management are viewed negatively in analysts' analyses of business outcomes relating to culture. Moreover, we show that the top three reasons analysts discuss culture are: 1) financial performance, valuation, and competitive advantage; 2) strategic alignment and execution; and 3) human capital management.

How do analysts' views of corporate culture compare with existing research on culture? Using values presented on websites of S&P 500 firms, Guiso, Sapienza, and Zingales (2015) find that those proclaimed values appear irrelevant. Yet, when employees perceive top management to be trustworthy and ethical, their firm performance is stronger. Using employee reviews to measure culture, both Guiso, Sapienza, and Zingales (2015) and Grennan (2019) conclude that traditional measures of corporate governance do not seem to have a significant impact on culture; Li, Chen, and Shen (2024) conclude that CEOs ultimately have moderate impact on corporate culture. Using job postings to measure corporate culture, Pacelli, Shi, and Zou (2022) establish a positive association between firms with a strong culture and their ability to attract and retain talent. After surveying 1,348 North American executives, Graham et al. (2022a) find the top three-ranked cultural values are results-oriented culture, community, and collaboration; the top three factors determining a firm's current culture are current CEO, market place, and owners; and the top three business outcomes affected by corporate culture are productivity, profitability, and creativity. Clearly, analysts' research on corporate culture offers insights that are both richer (i.e., more cultural values as well as their causes and effects identified) than and distinct from those derived from corporate websites, employee reviews, job postings, and/or surveys and interviews with corporate executives (e.g., according to analysts, the top three cultural values prevailing in modern corporations are innovation, adaptability, and customer-orientation, whereas according to executives, they are results-orientation, community, and collaboration).

To gain a better understanding of why analysts are interested in corporate culture, we employ regression analysis relating firm and analyst characteristics to the likelihood of analysts' featuring culture discussions in their reports, the number of values discussed, and their respective tones when discussing different cultural values. We find that firm size, sales growth, profitability, and major events such as top management turnover and deal-making are positively and significantly associated with, whereas leverage, earnings volatility, tangibility, ownership by large shareholders, and board independence are negatively and significantly associated with, analysts' discussing culture in their reports. Moreover, we find positive and significant associations between both executives' discussing culture in calls and the number of employees rating corporate culture on Glassdoor and analysts' featuring culture in their reports. Our findings on the negative influences of large shareholders and the positive influences of key corporate events (such as management turnover and M&As) are largely consistent with prior literature (e.g., Guiso, Sapienza, and Zingales 2015; Li et al. 2021). Our findings on the positive association between management discussing culture in calls (employees' posting culture-related reviews on Glassdoor) and analysts' featuring culture in their reports, and on the negative association between management discussing culture in calls and analysts' tones about culture in their reports are new, suggesting that analysts do pay attention to stakeholders and value-relevant intangibles such as corporate culture in their research.

In terms of analyst characteristics, we show that analysts who are women, are more experienced, and are affiliated with large brokers are more likely to discuss culture in their reports. The associations between certain analyst characteristics – gender, experience, and broker prestige – and coverage of culture help assuage concerns about analysts' indifference to or lack of insights into corporate culture.

Our cumulative evidence above suggests that equity analysts do have unique and significant insights into corporate culture from outside a corporation compared to insights gleaned largely from stakeholders within a corporation.

Finally, we explore whether there is any relationship between corporate culture and price formation. We find that analysts' positive tones in discussing culture are positively and significantly associated with their stock recommendations and target prices. In terms of economic significance, a change in tones from neutral to positive is associated with a 2.7 percentage point-increase in the probability of analysts upgrading their recommendations and a 1.4 percentage point-increase in target price forecast relative to the mean. Importantly, we also find that investors react positively and significantly to the positive tone in report text on culture, controlling for a report's quantitative and qualitative characteristics and analyst/firm characteristics. In terms of economic significance, a change in tones from neutral to positive results in an additional three-day abnormal return of 30.6 basis points around the release date of a report, corresponding to an $88.4 million increase in market value for an average firm in the sample. We conclude that analysts' research on culture offers new insights into its causes and effects, and that we are one step closer than prior work to establishing the culture-firm value link.

We contribute to the literature in three ways. First, our study provides new insights into corporate culture from the vantage points of important information intermediaries in capital markets – equity analysts. As such, our paper extends prior work that studies corporate culture using websites, employee reviews, surveys and interviews with executives, earnings conference calls, and/or job postings (see, for example, Guiso, Sapienza, and Zingales 2015; Grennan 2019; Au, Dong, and Tremblay 2021; Li et al. 2021; Briscoe-Tran 2022; Graham et al. 2022a, 2022b; Pacelli, Shi, and Zou 2022; Li, Chen, and Shen 2024). By applying generative AI to rich and granular textual data for information extraction, we reveal the causes and effects of different cultural values that have been difficult if not impossible to uncover in prior work. We further show that analysts' views of corporate culture differ from those of executives and employees. Our novel findings from capital market professionals' perspectives facilitate a better understanding of the mechanisms through which culture affects business outcomes, and have a wide range of management implications.

Second, our study contributes to the literature on big data and machine learning in finance, accounting, and economics (Gentzkow, Kelly, and Taddy 2019; Goldstein, Spatt, and Ye 2021). Contemporaneous research demonstrates the power of generative AI models in enabling capital market participants to derive new information and glean valuable insights from large quantities of textual data (Bai et al. 2023; Bybee 2023; Jha, Qian, Weber, and Yang 2023; Kim, Muhn, and Nikolaev 2023; Lopez-Lira and Tang 2023; Li, Tu, and Zhou 2023). Several papers explore the impact of generative AI on firm value and stock returns (Bertomeu, Lin, Liu, and Ni 2023; Eisfeldt, Schubert, and Zhang 2023; Babina, Fedyk, He, and Hodson 2024). Our study differs from current research in two important ways. First, we apply generative AI for complicated information extraction rather than text classification or prediction tasks, and demonstrate its effectiveness. Our approach yields a multi-faceted output that encompasses different cultural values and their respective causes or effects. Second, we combine the versatility of generative AI models such as ChatGPT, which are constrained by such factors as speed, cost, and context length limitations, with the efficiency of smaller large language models such as Bidirectional Encoder Representation from Transformers (BERT) models (Devlin, Chang, Lee, and Toutanova 2018), to effectively filter and retrieve the most relevant information.

On that note, our study highlights some unique considerations and design elements that must be addressed in order to harness the full potential of generative AI models in the context of financial text analysis. For instance, a step-by-step, chain-of-thought prompting strategy is beneficial for extracting perceived cause-effect relations – a task that requires high-level reasoning (Wei et al. 2022a). Furthermore, all generative AI models have intrinsic context length limitations, making it impossible to ask them to analyze a report in its entirety.[1] We demonstrate that feeding smaller segments related to corporate culture, while allowing for the dynamic augmentation of input segments by searching for relevant information from a full report, can enhance the overall capability of these models.

Third, our study contributes to the literature on equity analysts, particularly the strand applying textual analysis to research reports in order to gain insights into their information discovery and interpretation roles (e.g., Asquith, Mikhail, and Au 2005; Twedt and Rees 2012; Huang, Zang, and Zheng 2014; Huang et al. 2018; Bellstam, Bhagat, and Cookson 2021).[2] By applying generative AI to one of the largest report samples available, our research sheds light on the "black box" of analysts' fundamental research by not only underscoring culture as an integral input, but also elucidating the deductive processes analysts employ to transform qualitative, soft information into actionable insights (e.g., stock recommendations). As such, our big data-based research complements case study/interview/survey approaches

[1] For example, the combined input and output token length cannot exceed 4,000 for GPT-3.5-turbo, or about 3,000 English words. In addition, although some models can theoretically process longer inputs, their performance degrades when accessing and analyzing relevant information in the middle of a longer input (Liu et al. 2023).

[2] Using a sample of 1,126 reports by 56 All-America analysts over the period 1997-1999, Asquith, Mikhail, and Au (2005) construct a measure for the strength of arguments and show that such measure reduces, and sometimes eliminates the significance of the information available in earnings forecast or recommendation revisions. Using a sample of 2,057 reports in 2006 and a dictionary approach to measure tones in reports, Twedt and Rees (2012) find that the tone in a report contains significant information incremental to its quantitative content (e.g., earnings forecasts). Using a sample of 363,952 reports issued for S&P 500 firms over the period 1995-2008, Huang, Zang, and Zheng (2014) show that investors react more strongly to analyst reports that emphasize non-financial topics more than financial topics. Applying topic modeling, Huang et al. (2018) find that analysts both provide new information and interpret information released by corporate managers in their reports. Using a similar methodology and 665,714 analyst reports issued for S&P 500 firms over the period 1990-2012, Bellstam, Bhagat, and Cookson (2021) show that their textual-based measure of corporate innovation is more comprehensive and accurate than standard metrics of innovation output through multiple validation tests. Yet little is known about analysts' deductive process to produce their research output.

(e.g., Soltes 2014; Brown et al. 2015; Chi, Hwang, and Zheng 2023) to gain a better understanding of analysts' information production process.[3]

2. Overview of Generative AI for Cause-Effect Relation Extraction

2.1. Why generative AI?

Generative AI models such as ChatGPT and PaLM are gaining increasing popularity in social sciences, owing to their unprecedented performance in various applications. As a major advancement in the domain of natural language processing (NLP), these models exhibit attributes that are fundamentally different from conventional NLP methods (such as Latent Dirichlet allocation (LDA) (Blei, Ng, and Jordan 2003) or naïve Bayes models (Antweiler and Frank 2004; Li 2010)). At the heart of these models lies the transformer architecture, a form of neural networks that leverage self-attention mechanisms to weigh the influence of context words on the interpretation of a given word within a sentence (Vaswani et al. 2017). This architecture allows these models to capture long-range dependencies between words and understand contextual information effectively. During training, these models learn to predict missing words in a sentence, utilizing a vast amount of textual data to refine their predictions. This training method, termed self-supervised learning, equips the models with a broad understanding of human language. In addition to the structure and meaning of language (syntax and semantics), these models also learn general knowledge about the world and human experience. This "world knowledge" includes facts, concepts, and relationships that are commonly known or understood by humans.

[3] After examining a set of proprietary records from a large-cap firm, Soltes (2014) concludes that analysts' private interactions with management help them interpret firm news and better understand a firm's operations. Based on surveys and interviews of analysts, Brown et al. (2015) find that analysts' private communications with management are more helpful to their earnings forecasts and stock recommendations than their own primary research, recent earnings performance, and 10-K/10-Q filings, suggesting that analysts may be better-positioned to assess culture prevailing in an organization. Chi, Hwang, and Zheng (2023) show that analysts utilizing alternative data (such as job postings and employee reviews) produce more accurate forecasts, and their forecasts generate greater stock market reactions.

What do these transformer-based generative AI models offer that are different from conventional NLP methods? First, generative AI is known to develop emergent abilities as its model size scales up (Wei et al. 2022b). The models employed in our study have a substantially larger number of parameters in comparison to BERT models (Devlin et al. 2018),[4] effectively enabling these large language models (LLMs) to function as reasoning agents capable of comprehending and extracting complex information from analyst reports.[5] Two elements of these emergent abilities are zero-shot and few-shot learning, i.e., these LLMs are able to perform tasks, such as cause-effect relation extraction from analyst reports, on which they have not been explicitly trained and/or on which they have access to only very limited labeled data.

Second, we incorporate two new computational techniques on top of a generative AI model. The first, chain-of-thought prompting (Wei et al. 2022a) allows the model to break down a complex task into manageable subtasks. The second, retrieval augmented generation (Lewis et al. 2020) allows the model to dynamically search and retrieve additional information from a document to carry out the analysis when it recognizes that the context available falls short. Both techniques help strengthen the model's ability to handle complex tasks involving incomplete or fragmented context. To the best of our knowledge, we are the first to employ both techniques in finance or accounting applications.

2.2. Cause-effect relation extraction

---

[4] GPT-3.5-turbo is reported to have 175 billion parameters, and PaLM 2 has 540 billion parameters. In comparison, $BERT_{base}$ has 110 million parameters.

[5] Built on the transformer architecture, the BERT model and its variants, such as FinBERT (Huang, Wang, and Yang 2023), focus on predicting randomly masked missing words in text by taking into account the words that appear both before and after. Learning such bidirectional contexts makes it ideal for text classification tasks (e.g., sentiment analysis) and retrieval tasks (e.g., searching for relevant documents) (Reimers and Gurevych 2019). By contrast, GPT/PaLM models focus on predicting the next word based on the preceding context and are fine-tuned to follow a broad class of written instructions (Ouyang et al. 2022), which makes them more suited for text generation following human instructions, as exemplified in our study, in which we employ generative AI to extract cause-effect relations through specified prompts.

Transforming unstructured natural language (e.g., report text) into output that organizes information into different categories and relations requires an Information Extraction (IE) system. An IE system is designed to identify and understand entities and their relations from text (Etzioni et al. 2008). Within IE, Relation Extraction (RE) is a core task. RE generates triples such as (*'company A,' 'acquired,' 'company B'*), where "company A" and "company B" are entities, and "acquired" represents the relation between them. The relations derived from RE hold the potential for a wide range of applications. For example, the insights gained can aid theory development by generating new hypotheses (Mihăilă et al. 2013), and can serve as foundations for precise knowledge tasks such as fact-checking, question answering (Oh et al. 2013), and power prediction models in finance and economics (Radinsky et al. 2012).

Several of our research questions—identifying the cultural values common to corporations, and determining how such values are shaped and how they in turn influence business outcomes—align well with a form of IE known as Open Information Extraction (Open IE), in which the entities and relations are not pre-determined (Etzioni et al. 2008). In other words, we do not make any assumptions about predefined "bins" or categories into which different cultural values (e.g., innovation and teamwork), along with their causes and effects, can be placed. From a machine learning perspective, the lack of predefined categories renders supervised learning unsuitable for our open-ended research questions.

To this end, our paper employs generative AI for cause-effect relation extraction, thereby introducing a novel framework for analyzing financial text. Under this framework, we parse and represent analysts' analyses of corporate culture in a series of triples, each capturing an interaction between a cultural value (e.g., adaptability) and a business outcome (e.g., profitability). The nature of these triples is kept flexible, undergoing canonicalization only after the extraction. In our context, we impose two conditions on those triples: 1) the extracted relation is either a cause or an effect; and 2) one of the entities is a cultural value.

3. Extracting Analysts' Views of Corporate Culture from Reports

Figure 1 presents a flowchart of how we apply generative AI to extract analysts' views of corporate culture from their research reports.

3.1. Identifying corporate culture-related segments in reports

We download from Thomson One's Investext database 2,434,782 reports covering S&P 1500 constituent firms over the period 2000–2020. We convert these reports from PDF format to plain text. An inherent challenge of this step is the loss of paragraph structure in reports. Moreover, even if the structure could have been maintained, differentiating between headers, bullet points, and coherent paragraphs in a report is not straightforward.

To address these issues, we employ a common text segmentation technique developed by Choi (2000), the C99 algorithm. This method coalesces individual sentences into larger, more meaningful *segments* that align more closely with coherent thoughts or ideas in the text. Consequently, we use these segmented units, rather than individual sentences, as the unit of analysis for our study.

Next, we implement a machine learning model to filter out boilerplate content from reports. The data set for training consists of segments in reports produced by the top 20 brokers, with positive examples identified as the most frequently repeated segments and negative examples as those least repeated. We fine-tune a BERT model to classify boilerplate segments automatically, and the trained model demonstrates high accuracy. Table IA1 in the Internet Appendix lists predicted boilerplate probabilities and boilerplate examples, sorted by decile. We retain segments with a boilerplate probability of 0.22 (the sample median) or lower.

We then identify culture-related segments in reports through a two-step procedure.[6]

In step 1, we start with an exhaustive text search using two sets of keywords. The initial set of keywords is based on the word set explicitly about corporate culture, identifying a total of 5,541 relevant segments.[7] We also employ a second, more flexible set of keywords, which matches all segments containing the word "culture(s)" or "cultural," excluding those already identified to avoid duplication. This second search results in a larger set of 46,795 segments. Some segments, however, are in the biological or societal context. We apply generative AI for word sense disambiguation (WSD) to filter out irrelevant segments. Table IA2 Panel A in the Internet Appendix shows the prompt used. After filtering, this step results in a total of 41,038 relevant segments.

In step 2, we fine-tune a BERT model to identify culture-related segments that lack specific keywords. The construction of our training set involves using segments, identified in step 1 as containing relevant keywords, as positive examples (culture = 1). Conversely, we include randomly selected segments without those keywords as negative examples (culture = 0). This training set is used to fine-tune the model, which is then deployed across all segments (excluding those identified in step 1). We sort segments by their predicted probabilities of relating to culture, and focus on the top 5% segments with the highest predicted probabilities. To improve the filtering, we use generative AI to further scrutinize these segments and retain only segments that are explicitly about corporate or organizational culture. Table IA2 Panel B in the Internet Appendix shows the prompt used. This step adds

---

[6] Li et al. (2021) develop cross-validated measures of corporate culture by applying word embedding to earnings conference call transcripts; their measures offer broad coverage of firms over time. In contrast, our paper aims to provide insights into analysts' views of cultural values in terms of their perceived causes and effects, which is a much harder task than measuring culture. In pursuit of our goal, we emphasize accuracy in identifying culture-related segments in analyst reports over coverage, using two word lists and a machine learning model as discussed in this section.

[7] We use the following phrases for exact matching: "corporate culture," "company culture," "company's culture," "firm culture," "firm's culture," "organizational culture," "workplace culture," "business culture," and "culture in the company."

51,434 segments. Our final data set comprises 92,472 culture-related segments in reports (41,038 segments from step 1 + 51,434 segments from step 2).[8]

Figure 2 plots the intensity of culture-related segments in analyst reports over time and by report section over the period 2000–2020. The horizontal axis indicates report year, and the vertical axis indicates report section, binned into 20 equal sections from the start to the end of a report. The color gradient depicts the intensity of culture-related segments, computed as the number of these segments normalized by the number of reports in a year (we multiply this variable by a hundred).

We make two observations. First, there is a shift in the location of culture-related segments in reports. At the beginning of the sample period, these segments are generally scattered throughout a report. In more recent years, they appear more concentrated in the first half of a report. This shift suggests that, over time, analysts are more aware of the importance of culture, and hence they position their culture-related analyses in the front end of their reports. Second, there is a marked rise in the intensity of culture-related segments in reports in recent years following a discernable dip during and in the aftermath of the Great Recession (2008–2012); this dip could potentially be attributed to analysts' heightened focus on financial performance and cost-cutting in a period of economic uncertainty.

3.2. Cause-effect relation extraction using generative AI

In this section, we combine generative AI with chain-of-thought (CoT) prompting and retrieval augmented generation (RAG) to perform RE from analyst reports.

CoT prompting is a reasoning methodology that aids generative AI in executing complex tasks. It operates by instructing the model to perform step-by-step reasoning or

[8] It is informative to compare the number of culture-related segments in reports with the number of segments mentioning other value-relevant events or economic condition. For example, using the phrases "joint ventures," "joint venture," or "JV," we get 124,633 segments; using the phrases "strategic alliances," or "strategic alliance," we get 7,578 segments; for M&A-related mentions, we get 1,072,022 segments; and for "inflation," we get 229,490 segments.

problem-solving, mimicking the progressive nature of the human thought process. When applying the CoT prompting, we break down the elaborate task of RE into smaller, manageable steps. This simplification not only improves the performance of the model, but also provides interpretability into the model's problem-solving process, allowing for close inspection of each stage of reasoning and offering insights into how the model reaches its conclusions.

In the baseline case, when applying the CoT prompting, we direct the model to extract cultural values along with their causes and effects from each segment. The model is then asked to construct cause-effect relation triples that encapsulate the extracted information in a standardized format. It is possible for a single segment to have multiple cause-effect relations or none. Table 1 Panel A shows the CoT prompt used in our analysis.

In some cases, the segment available might be too concise, depriving the model of the necessary context to extract a causal relation. In such cases, we turn to RAG, which dynamically integrates external information from relevant segments into the problem-solving process. Specifically, if the model outputs "I need more context" during the CoT prompting, we perform a semantic search to retrieve the top five most relevant segments in the same report, plus those segments that are immediately before or after the focal segment, that could provide additional context for the analysis.[9] Table 1 Panel B shows the prompt used. The model is then provided with both the focal segment and the additional context retrieved via a semantic search to implement RE. Should no causal relation be extracted even with the additional context, the model will produce an empty output in the relevant field.

There are two main concerns when applying generative AI: look-ahead bias and hallucination. Look-ahead bias refers to using future information not available at the time of prediction. Our application focuses on extracting analysts' views of corporate culture rather

[9] When additional segments exceed the token limit of the model, we select those that are most similar to the focal segment within the limit.

than making out-of-sample predictions. Moreover, our regression analyses primarily use the extracted cultural values based on information available at the time of a report's publication. In short, our approach and research design by construction mitigate look-ahead bias concerns. Hallucination refers to generated text that is unfaithful or ignores source material (Maynez, Narayan, Bohnet, and McDonald 2020). A key strategy to mitigate hallucination is to improve the alignment between the input and the generated output (Ji et al. 2023). Our approach ensures that the input segments align closely with what the prompt asks for because the segments either contain explicit culture-related keywords or are selected through a multi-stage filtering process. Moreover, our CoT prompting guides generative AI to strictly reason within a report's content in a step-by-step manner. Finally, we also instruct generative AI to output "N/A" if pertinent information is absent, further reducing hallucination risk. Human annotation in Section 3.4 will verify the model's adherence to the context and knowledge present in reports.

In summary, combining generative AI with CoT prompting and RAG allows us to leverage the strengths of both approaches – structured problem-solving and dynamic contextual integration – to effectively extract cause-effect relations involving corporate culture from analyst reports.

3.3. Canonicalization of extracted information

After extracting cultural values from analyst reports using generative AI, we conduct *canonicalization*, a process that converts information with multiple representations/forms into a standard or canonical format. The initial output from the model has more than 15,100 variations of phrases referring to corporate culture, which we consolidate in two stages.

In the first stage, whenever possible, we manually examine and categorize 881 culture-related phrases that show up ten or more times in reports into broad cultural values guided by prior work listed in Table IA3 in the Internet Appendix.[10] Individually, we categorize those phrases. When any disagreement arises, we turn to the original reports to reach a consensus. We end up with ten cultural values defined as follows:

- **Adaptability**: Cultures emphasizing cultural shifts, adaptations, transformations, or resistance to change. For example, "cultural change," "cultural transformation," "culture reset," and "insular culture."

- **Customer-oriented**: Cultures prioritizing sales or customer relations. For example, "sales culture," "customer-centric," and "brand-focused."

- **Innovation**: Cultures emphasizing innovation, growth, and entrepreneurship. For example, "entrepreneurial," "innovation," "growth-oriented," and "technology-driven."

- **Integrity**: Cultures valuing ethical behavior, fair practices, and community relationships. For example, "enlightened hospitality," "ESG-focused," "fair business practices," and "blame culture."

- **Operations-oriented**: Cultures underscoring efficiency, productivity, and cost control. For example, "decentralized," "efficiency-focused," and "cost-conscious."

- **People-oriented**: Cultures centering on employee well-being, diversity, inclusion, empowerment, and talent growth. For example, "employee-centric," "diverse and inclusive," and "talent-focused."

- **Results-oriented**: Cultures focusing on performance, competitiveness, and results. For example, "competitive," "performance-driven," and "pay-for-performance."

- **Risk control**: Cultures stressing risk management, credit, and financial prudence. For example, "conservative credit culture," "risk-averse," and "culture of compliance."

- **Teamwork**: Cultures highlighting collaboration, integration, and team-orientation. For example, "collaborative," "integration," "cultural fit," and "partnership-oriented."

- **Miscellaneous**: This category contains various unspecific and less-frequently occurring cultural values, or those that do not easily fit into the above categories. For example, "positive culture" and "distinct culture."

[10] We are mindful that cultural values discussed by analysts in their reports might differ from those espoused by firms or their management, or experienced by employees, as examined in prior work. Given analysts' role as an information intermediary, assessing whether their views do differ is an important research question this paper addresses.

In the second stage, we employ generative AI to categorize culture-related phrases that show up fewer than ten times in reports. Table 1 Panel C shows the prompt used. To ensure that the examples provided in the prompt are relevant to the phrase being analyzed, we use a dynamic approach based on a nearest-neighbor criterion. Specifically, we embed each phrase using a model based on the BERT architecture, then compute the cosine similarity between a focal phrase (to be categorized) and each of the 881 phrases that we have manually categorized. Based on the cosine similarity score, we supply the most similar (already categorized) phrase and its corresponding categorized cultural value to the focal phrase as an example in the prompt.

We employ similar techniques to categorize the causes or effects of each cultural value. The process again starts with the use of a BERT-based model to embed all identified causes and effects into fixed-size vectors. We then cluster these vectors into 50 groups using agglomerative hierarchical clustering. The clusters provide an initial understanding of the common patterns and relationships among the wide array of causes/effects. After clustering, we manually inspect the clusters and combine them into major categories of causes and effects. Each major category is supported by five to ten human-annotated examples. After the major categories are established, we employ generative AI to classify the extracted causes and effects into those major categories with dynamically generated examples in the prompt. Table 1 Panel D shows the prompt used. Table IA4 in the Internet Appendix provides some representative examples of the extracted cultural values and their causes and effects.

Finally, we conduct canonicalization of numerous variations of cause-effect relation triples. Consider an example (*'aggressive sales culture,' 'resulted in,' 'massive scandal'*). The canonicalization process first assigns one "end" of the triple to one of the cultural values discussed above using the mapping created in the prior step. In this case, the phrase "aggressive sales culture" is assigned to the cultural value "Customer-oriented." We then use generative AI to classify the nature of the relationship into "->", "<-" or "<->" wherein the arrow indicates the direction of the cause or effect. Table 1 Panel E shows the prompt used in the analysis. For example, the phrase "*provides opportunity for*" is canonicalized as ->, "*threatened by*" as <-, and "*align with*" as <->. Panel F shows the prompt to canonicalize reasons for analysts to discuss culture.

A more detailed description of our generative AI-based method is provided in the Internet Appendix. Table IA5 in the Internet Appendix provides examples of snippets in culture-related segments, and of extracted cultural values and their respective causes and effects.

3.4. Model performance

In this section, we evaluate the performance of three generative AI models. Two of these models belong to the ChatGPT model series developed by OpenAI: *GPT-4-1106-preview* (hereafter referred to as GPT-4) and *GPT-3.5-turbo-1106* (referred to as GPT-3.5). The third model belongs to Google's Vertex AI PaLM 2 model series, specifically *text-bison-001* (referred to as PaLM).[11] Our evaluation involves manually annotating cultural values, causes, effects (and thereby cause-effect relations), and tones in 200 randomly selected culture-related segments, and comparing them with the relations and tones extracted by generative AI models. Table 2 presents the model performance results using accuracy, precision, recall, and F1 scores.

In terms of accuracy, GPT-4 achieves 94.5%, 91.5%, 94.0%, and 80.5% for cultural values, causes, effects, and cause-effect relations, respectively. GPT-3.5 achieves 91.0%, 86.5%, 88.5%, and 71.0%, respectively. PaLM achieves 81.5%, 80.0%, 85.5%, and 67.0%, respectively. The accuracy for cause-effect relations is lower than that for other tasks because

[11] To ensure reproducibility of our findings, we set temperature to 0 and a deterministic random seed (when applicable) in all API calls. We keep other hyperparameters at their default values.

a false extraction of any of the three elements (cultural values, causes, and effects) will be considered a false case (false positive/negative). We find similar results using precision, recall, and F1.

In terms of tones in culture-related segments, we show that GPT-4 achieves an accuracy of 93.0%, outperforming GPT-3.5 and PaLM, which achieve an accuracy of 91.5% and 88.5%, respectively. As a comparison, we also examine the performance of FinBERT developed by Huang, Wang, and Yang (2023) to capture tones in culture-related segments. We find that FinBERT only achieves an accuracy of 72.5%, primarily because it captures the overall tone of a segment, whereas generative AI (like humans) is capable of detecting the tone specific to the culture-related text within a segment. Our findings echo Lopez-Lira and Tang's (2023) conclusion that generative AI with its reasoning capabilities outperforms other machine learning models in capturing tones.

In summary, we find all three models exhibit good model performance when extracting cultural values, causes, effects, cause-effect relations, and tones in textual data. GPT-4 has a slight edge compared to GPT-3.5 and PaLM. However, the disparity in performance among these models is relatively small, while GPT-4 operates at a significantly slower speed and higher cost than GPT-3.5. In the remainder of this paper, we employ GPT-3.5 as our primary AI model to conduct analyses.

4. Revealing Analysts' Views of Corporate Culture from Reports

4.1. Overview of analysts' views of corporate culture

Table 3 provides the frequency count and share of culture-related segments relating to individual cultural value, cause, and effect in descending order. According to analysts, the top three cultural values are innovation (16.9%), adaptability (15.9%), and customer-oriented culture (11.5%), and the bottom three cultural values are integrity (3.7%), teamwork (5.3%), and people-oriented (6.5%). The top three drivers of cultural changes are business strategy (18.4%), management team (12.4%), and strategic transformation (8.8%), and the bottom three drivers are the COVID-19 pandemic (0.1%), workplace safety (1.0%), and regulatory issues (1.3%). The top three business outcomes that culture shapes are market share and growth (16.2%), profitability (13.6%), and employee satisfaction (8.4%), and the bottom three outcomes are misconduct (0.5%), internal conflicts (1.0%), and diversity, equity, and inclusion (1.3%).

We compare our findings with the survey data of 1,348 North American executives from mostly private firms (about half of them are family firms) reported by Graham et al. (2022a). According to the survey, executives view the following words/phrases as best describing the current culture at their firms: results-oriented (48%), community (39%), and collaboration (32%). When executives were given a list of terms related to business ethics (e.g., compliance), innovation (e.g., creativity), and productivity (e.g., profitability), and were asked to what extent their firms' cultures were affected by these terms, the top three factors influencing a firm's current culture as chosen by the executives were current CEO (55%), market place (35%), and owners (32%). Clearly, analysts' views of corporate culture are different from what we learn about culture based on surveys of executives. Huang et al. (2018) compare analyst reports and earnings calls and find that about a third of the reports discuss (new) topics not mentioned in calls. Our comparison of analysts' and executives' views of corporate culture is consistent with their findings.

Figure IA1 in the Internet Appendix provides an overview of the temporal trends in culture-related segments in analyst reports. In Panel A, we show a clear dip in such coverage around the Financial Crisis of 2008–2010. In Panel B, we separate culture-related segments by cultural value. We note that, over time, adaptability, innovation, and people-oriented (customer-oriented and results-oriented) values are gaining (losing) importance.

Figure IA2 provides an overview of the Fama-French 12 industry distributions of culture-related segments in analyst reports. In Panel A, we show that finance, shops (i.e., wholesale and retail), and business equipment are the three industries whose analysts pay the most attention to corporate culture, whereas utilities, telecom, and consumer durables are the three industries whose analysts pay the least attention. Consistent with our intuition, we show that analysts tend to focus on customer-oriented culture in the wholesale and retail industry, and risk control culture in the utilities and finance industries.

Figure IA3 provides a breakdown of culture-related segments by cultural value and tone. We show that, in general, analysts are less likely to use negative tones when analyzing culture. Among all cultural values, analysts are slightly more negative when writing about risk control, adaptability, and integrity.

4.2. The cause-effect knowledge graph

Importantly, our generative AI-fueled method allows us to identify the events, people, and/or systems that significantly influence a specific cultural value, and to determine which cultural values are most impactful for various business outcomes. Figure 3 plots the cause-effect knowledge graph capturing analysts' views of corporate culture.

The graph is divided into three columns of entities: 1) the left column lists the seventeen drivers of culture grouped by events, people, and systems suggested by Guiso, Sapienza, and Zingales (2015), Graham et al. (2022a, 2022b), and Grennan and Li (2023) (omitting the miscellaneous category); 2) the center column lists the nine cultural values (omitting the miscellaneous category); and 3) the right column lists the sixteen effects of culture (omitting the miscellaneous category). The height of each entity (a cultural value, a driver, or an effect) denotes the number of relevant segments, thereby providing an intuitive visual representation of each entity's prominence. The width of each link denotes the number of segments mentioning a particular cause or effect relation. For an uncluttered depiction, we only retain the top two most frequent links for each entity in Figure 3.

Focusing on the links coming out of the cause column (on the left side of Figure 3), analysts identify that business strategy is one of the top two influencing factors for a large number of cultural values – customer-oriented, operations-oriented, innovation, results-oriented, adaptability, integrity, and risk control, and management team for teamwork, innovation, results-oriented, adaptability, integrity, risk control, and people-oriented values. In contrast, strategic transformation is one of the top two influencing factors for operations-oriented and adaptability. Focusing on the links coming into the value column (in the middle of Figure 3), analysts identify that M&As (an event) and management team (people) are the top two factors shaping the cultural value of teamwork, and business strategy (a system) and management team (people) are the top two factors shaping the cultural value of integrity. Focusing on the links coming out of the value column (in the middle of Figure 3), analysts identify the cultural values of innovation and adaptability as having an impact on almost all aspects of business operations, ranging from market share and growth to corporate environmental, social, and governance (ESG) practices, while other values such as customer-oriented, operation-oriented, or people-oriented have less impact. For example, the cultural value of customer-oriented is one of the top two values shaping business outcomes, including customer satisfaction and market share and growth, and people-oriented is one of the top two values shaping employee satisfaction and diversity, equity, and inclusion. Focusing on the links coming into the effect column (on the right side of Figure 3), analysts identify customer-oriented and innovation as the top two values shaping market share and growth, and operations-oriented and results-oriented as the top two values shaping profitability. Notably, the above discussions are merely the proverbial tip of the iceberg, with many additional granular insights into the complex relations among the nine cultural values, seventeen causes, and sixteen effects available from these links. To the best of our knowledge, such insights are wholly absent from the existing literature on corporate culture.

In the effect column, we utilize a color-coding scheme to represent tones in culture-related segments, with darker shades signifying more negative tones. We find that innovation, customer satisfaction, and employee satisfaction are viewed positively, whereas misconduct, internal conflicts, and risk management are viewed negatively in analysts' analyses of business outcomes relating to culture. In our sample, 66.7% of analysts have positive tones when discussing culture. By comparison, Graham et al. (2022a) note that 63% of the surveyed executives express positive sentiments when discussing culture.

Lastly, we ask the model to list reasons for analysts' discussing culture and find that the top three reasons are: 1) financial performance, valuation, and competitive advantage; 2) strategic alignment and execution; and 3) human capital management.

Overall, Figure 3 provides a comprehensive visualization of analysts' views of corporate culture spanning different cultural values, their causes, and their effects, which differs from corporate insiders' perceptions of cultural values and their roles in business outcomes (see, for example, Guiso, Sapienza, and Zingales 2015; Grennan 2019; Graham et al. 2022a, 2022b). Knowledge gained from revealing analysts' views of culture can serve as fact-checking of research ideas and provide a roadmap for empirical research on culture.

5. Understanding Analysts' Views of Corporate Culture

Unlike values or quantities gleaned from financial statements, the notion of corporate culture is somewhat nebulous and thus raises a number of questions regarding our research premise: Do sell-side equity analysts possess value-relevant insights into corporate culture? Moreover, is it possible that analysts are simply reiterating management's narratives in earnings conference calls, with little originality in their discussions of culture? We address these concerns in this section. We also conduct regression analysis to gain a better understanding of the determinants of analysts' featuring culture in their reports.

5.1. Sample formation and overview

We download from Thomson One's Investext 2,434,782 reports covering S&P 1500 constituent firms over the period 2000–2020. We obtain report date, gvkey, lead analyst name (including last name and first name initial), and broker name from the meta file. Section B of the Internet Appendix describes how we match analyst name in a report to analyst ID in the Institutional Brokers Estimates System (I/B/E/S) database, in order to construct analyst characteristic variables. Table 4 lists the steps taken and filters applied to form our main firm sample, comprising 28,860 firm-year observations representing 2,569 unique firms. Our firm-analyst-year sample comprises 142,868 observations representing 2,484 firms covered by 4,043 analysts.

Table 5 provides the summary statistics for the different samples used in our analyses. All continuous variables are winsorized at the 1st and 99th percentiles, and the dollar values are in 2020 dollars. Variable definitions are provided in the Appendix. Panel A presents the summary statistics for the firm-year sample and firm-analyst-year sample. At the firm-year level, our measure of the frequency of analysts' discussions of culture, *Culture discussion*, has a mean of 0.425, suggesting that there is at least one analyst discussing culture in about 40 percent of firm-year observations. The average number of cultural values discussed is slightly over two. On average, analysts have a positive tone when discussing culture; the sample average tone is 0.676 (1 for positive tone, 0 for neutral, and -1 for negative tone). Among the nine cultural values, innovation and adaptability are the most frequently mentioned, whereas integrity and teamwork are the two least frequently mentioned, based on analyst reports. At the firm-analyst-year level, we show that about a tenth of firm-analyst-year observations discuss culture. Table IA6 in the Internet Appendix presents the Pearson correlation matrices of the firm-year sample and the firm-analyst-year sample. Examination of the correlation matrices suggests that multicollinearity is unlikely to be an issue.

5.2. Do analysts have insights into corporate culture?

To address the concern that analysts do not understand intangibles such as corporate culture, we note that analysts gain in-depth knowledge of the firms they follow through both formal and informal channels, including attending earnings conference calls, conducting site visits, and meeting with top and divisional managers (Soltes 2014; Brown et al. 2015). All of the above provide a window into an organization's culture. Moreover, prior work finds that analysts possess value-relevant non-financial information about the companies they cover (Previts et al. 1994; Huang, Zang, and Zheng 2014; Brown et al. 2015; Bellstam, Bhagat, and Cookson 2021); and corporate culture is part of firms' intangible assets. We further note that the nine cultural values featured in analyst reports are typically not exact copies of cultural values as viewed on corporate websites (Guiso, Sapienza, and Zingales 2015), discussed in calls (Li et al. 2021), experienced by employees (Grennan 2019), or mentioned by executives in surveys and interviews (Graham et al. 2022a, 2022b), which suggests, as noted above, that analysts possess unique insights into corporate culture (with the caveat that analysts' discussions of these cultural values in their reports may nonetheless occasionally repeat related talking points from earnings calls) (see Table IA3 in the Internet Appendix for a comparison). We address this concern by comparing analysts' views in a sample of firm-year observations that hold calls with those in a sample of firm-year observations that do not hold calls; we further compare analysts' views in a sample of firm-year observations in which executives discuss corporate culture in greater depth during calls with those in a sample of firm-year observations in which they do not do so. Table 6 presents the results.

In Panel A, we find no significant difference in analysts' views between the two samples across a number of summary measures and individual cultural values. In a number of cases in which views differ, analysts tend to pay significantly more attention to customer-oriented, results-oriented, and risk control cultures in firms that do not hold calls than analysts covering firms that do hold calls, suggesting that analysts are not simply repeating what they have heard in calls. In Panel B, we show that even for the sample of firm-year observations in which executives discuss the least about culture in calls, analysts who follow those firms still discuss more about risk control culture and have similar amounts of discussion of operations-oriented and results-oriented cultures as those in which executives talk more about culture.

In summary, our evidence suggests that analysts possess meaningful insights about corporate culture.

5.3. Firm characteristics and analysts' discussing corporate culture

To examine whether and how firm characteristics are related to the likelihood of analysts' research on culture in their reports, we employ the following regression specification at the firm-year level:

$$Y_{i,t} = \alpha + \beta \times Firm\ characteristics_{i,t-1} + Ind\ FE + Year\ FE + \varepsilon_{i,t}, \qquad (1)$$

where the dependent variables $Y_{i,t}$ are: whether analysts discuss culture (or not), the number of cultural values discussed, and their tones when discussing culture. Firm characteristics largely follow prior work (Guiso, Sapienza, and Zingales 2015; Li et al. 2021; Li, Chen, and Shen 2024). We include industry fixed effects to control for the effect of time-invariant industry factors on analysts' discussions of culture, and year fixed effects to control for changing trends in analysts' awareness of culture.

Table 7 Panel A presents the results when the dependent variable is whether a firm's analysts discuss culture or not. Columns (1)-(2) present the regression results including the basic set of firm-level controls. We find that firm size, sales growth, ROA, the number of management turnover cases, and the number of M&A deals are positively and significantly, whereas firm age, leverage, tangibility, earnings volatility, ownership by large shareholders (using 5% as the shareholding cutoff), and board independence are negatively and significantly, associated with the likelihood of analysts' discussing culture in their reports. In columns (3)-(4), we further control executives' and employees' discussions of culture. We note that both management discussing culture in calls and the number of employees rating culture and values on Glassdoor are positively and significantly associated with the likelihood of analysts' discussing culture in their reports.

Limiting to a subsample of firm-year observations with culture-related segments in analyst reports, we examine the determinants of analysts' discussing different values and their respective tones. Panel B presents the results. In columns (1)-(2), we show that firm size, sales growth, ROA, the number of management turnover cases, and the number of M&A deals are positively and significantly, whereas leverage and ownership by large shareholders are negatively and significantly, associated with the number of values discussed in analyst reports. In columns (3)-(4), we show positive and significant associations between both management discussing culture in calls and the number of employees rating culture and values on Glassdoor, and the number of values discussed by analysts. In terms of tones in analysts' discussions of culture in columns (5)-(8), we show that firm size, earnings volatility, the number of management turnover cases, and management discussing culture in calls are negatively and significantly, whereas ROA and tangibility are positively and significantly, associated with analysts' tones. Our findings on the negative influences of large shareholders and the positive influences of key corporate events (such as management turnover and M&As) are largely consistent with prior literature (e.g., Guiso, Sapienza, and Zingales 2015; Li et al. 2021). Moreover, our findings as noted above on the positive associations between management's, employees', and analysts' discussing corporate culture, and on the negative association between management's discussion of culture in calls and analysts' tones about culture in their reports are new, suggesting that analysts do have unique insights into corporate culture in their research.

5.4. Analyst characteristics and their discussing corporate culture

To examine what analyst characteristics are associated with the likelihood of their writing about culture in reports, we employ the following regression specification at the firm-analyst-year level:

$$Y_{i,j,t} = \alpha + \beta \times Analyst\ characteristics_{j,t-1} + Firm \times Year\ FE + Broker\ FE + \varepsilon_{i,j,t}, \tag{2}$$

where the dependent variables $Y_{i,j,t}$ are: whether an analyst discusses culture (or not), the number of cultural values discussed, and her tone in discussing culture. Our analyst characteristics largely follow prior literature (e.g., Clement and Tse 2005). In addition to analyst characteristics, we include firm × year and/or broker fixed effects to control for time-varying unobservable firm and/or broker characteristics that may affect analysts' coverage decisions and/or their decisions to discuss culture in reports. Table 7 Panel C presents the regression results.

Columns (1) and (4) present the regression results when the dependent variable is whether an analyst discusses culture or not. We show that analysts who are women; who have more general and firm-specific experiences; and who are affiliated with large brokers are each more likely to discuss corporate culture. Columns (2) and (5) present the regression results when the dependent variable is the number of values discussed by an analyst. We show that analysts with longer forecast horizons, more general experience, and more industry coverage are positively and significantly with the number of values discussed in their reports. Columns (3) and (6) present the regression results when the dependent variable is the tone of analysts' discussion of culture. We show little explanatory power from analyst characteristics considered.

We conclude that the significant associations between certain analyst characteristics – gender, experience, scope of coverage, and broker prestige – and their coverage of culture help assuage concerns about analysts' indifference to or lack of insights into corporate culture.

In summary, our analyses thus far suggest that generative AI models have the potential to reveal new insights into corporate culture from the vantage points of sell-side equity analysts, and that those insights differ from corporate insiders' views of culture. We next answer our final research question: What is the relationship between corporate culture and price formation?

6. Analysts' Views of Culture and Stock Price Implications

Prior studies show that analysts' fundamental research contributes to stock price formation (see, for example, Womack 1996; Brav and Lehavy 2003; Loh and Stulz 2011; Huang, Zang, and Zheng 2014; Kecskés, Michaely, and Womack 2017). In this section, we examine whether and how analysts' views of culture impact price formation. The analysis is at the report level. Section C of the Internet Appendix describes how we match reports from Investext in our sample to the I/B/E/S forecast data.

6.1. Analysts' views of corporate culture and their research output

To examine the relationship between analysts' views of culture and their research output, we focus on stock recommendations and target prices because both measures capture a firm's long-term prospects (e.g., Brav and Lehavy 2003; Loh and Stulz 2011), which aligns well with the role of corporate culture in long-term value creation (e.g., Guiso, Sapienza, and Zingales 2015; Li, Liu, Mai, and Zhang 2021). Moreover, a strong culture could over time boost a firm's cash flows and/or lower its discount rate; either or both outcomes would be reflected in stock recommendations and target prices. Given the importance of sentiment in textual data (see, for example, Antweiler and Frank 2004; Tetlock 2007; Loughran and McDonald 2011; Huang, Zang, and Zheng 2014), the key variable of interest is ChatGPT's classification of analysts' tones in culture-related segments, *Tone*. We also examine tones at more granular levels such as cultural value-specific tones (e.g., *Adaptability tone*) as well as control for tones in the remainder of a report, *Non-culture tone*. Table 8 Panel A presents the summary statistics for the key variables. We note that they are largely consistent with prior literature (e.g., Bradshaw, Brown, and Huang 2013; Huang, Zang, and Zheng 2014; Kecskés, Michaely, and Womack 2017). Panels B-D present the regression results. We include firm × year and analyst fixed effects to control for time-varying unobservable firm characteristics that may affect analysts' coverage decisions and analyst innate skill or preferences relating to their discussions of corporate culture, respectively.

In Panel B, we show that *Tone* is positively and significantly associated with stock recommendations and target prices, suggesting that analysts' sentiments on culture play a significant role in their stock recommendations and target price forecasts. In terms of economic significance, a change in *Tone* from negative to neutral (or from neutral to positive) is associated with a 2.7 percentage point-increase in the (linear) probability of analysts upgrading their recommendations and a 1.4 percentage point-increase in target price forecast relative to the mean.[12] For comparison, a change in *Non-culture Tone* from negative to neutral (or from neutral to positive) is associated with a 7.8 percentage point-increase in the (linear) probability of analysts upgrading their recommendations and a 5.8 percentage point-increase in target price forecast relative to the mean.[13] Given that the non-culture segments

[12] The 2.7 percentage point-increase is calculated from 1 × 0.109 × 100/4 where the denominator 4 is the range of stock recommendation from -2 to 2, and the 1.4 percentage point-increase is calculated from 1 × 1.653 × 100/118.141 where the denominator 118.141 is the sample average target price (%).

[13] The 7.8 percentage point-increase is calculated from 1 × 0.312 × 100/4, and the 5.8 percentage point-increase is calculated from 1 × 6.882 × 100/118.141.

regarding a firm's financial performance are more directly related to stock recommendations and target price forecasts, the effect of culture-related *Tone* is noteworthy.

In Panel C, we further break down analysts' views into their respective views on each of the ten cultural values (including the miscellaneous category), and examine whether and how these tone variables are associated with analysts' recommendations and target price forecasts. We show that all cultural value-related tone variables are positively and significantly associated with stock recommendations and target prices. In terms of economic importance, the top three drivers of stock recommendations are *Adaptability tone*, *Innovation tone*, and *Miscellaneous tone*, and the top three drivers of target price forecasts are *Innovation tone*, *Adaptability tone*, and *Teamwork tone*.

In Panel D, we employ a subsample of reports associated with recommendation (target price) revisions.[14] We show that there is a negative and significant association between *Tone* and recommendation downgrades (target price downward revisions), suggesting that analysts do take into account their views of culture when making revisions.

6.2. The information content of analysts' views of corporate culture

To investigate the information content of analysts' views of corporate culture in reports, we employ an event study relating three-day cumulative abnormal returns (CAR) around the report date, to measures of analysts' views of culture controlling quantitative and qualitative summary measures of a report, and analyst and firm characteristics (Huang, Zang, and Zheng 2014; Huang et al. 2018).[15] Table 9 presents the regression results.

[14] Following Huang, Zang, and Zheng (2014), a report is considered a revision if the report date is within the five-day window centered around the I/B/E/S recommendation (target price) announcement date when an analyst revises her recommendation (target price forecast). All other reports are considered reiteration reports.
[15] For this analysis, we remove 892,467 reports due to companies with multiple reports, 12,074 reports due to companies issuing earnings announcements, and 5,058 reports due to companies issuing earnings guidance, in the CAR window.

We show that *Tone* is positively and significantly associated with *CAR[-1,+1],* suggesting that culture discussions in a report provide information beyond that provided by its quantitative and qualitative measures. In terms of economic significance, a change in *Tone* from negative to neutral (or from neutral to positive) results in an additional three-day abnormal return of 30.6 basis points around the report date, corresponding to an $88.4 million increase in market value for an average firm in the sample.[16] For comparison, a change in *Non-culture tone* from negative to neutral (or from neutral to positive) results in an additional three-day abnormal return of 113.6 basis points, corresponding to a $328.3 million increase in market value for an average firm in the sample.[17] It is worth noting that the effect documented above is the direct information effect of analysts' views of culture in reports, and that there are also indirect effects via stock recommendation and target price revisions shown in Table 8.

We conclude that analysts incorporate their views of corporate culture in their research output, and that investors significantly react to report text on culture.

7. Conclusions

Our study is among the first in finance, accounting, and economics to apply generative AI models as reasoning agents on analyst reports to gain insights into analysts' views of corporate culture.

We employ generative AI (ChatGPT) to analyze 2.4 million analyst reports between 2000 and 2020. Generative AI organizes analysts' views into a knowledge graph that links different cultural values to their perceived causes and effects. In terms of influencing factors

---

[16] The 30.6 basis points increase is calculated from $1 \times 0.306 \times 100$, and the $88.4 million increase in market value of equity is calculated from $0.306\% \times \$28.9$ billion where $28.9 billion is the sample average market capitalization.

[17] The 113.6 basis points increase is calculated from $1 \times 1.136 \times 100$, and the $328.3 million increase in market value of equity is calculated from $1.136\% \times \$28.9$ billion.

for corporate culture, analysts identify business strategy as a key factor for a large number of cultural values – customer-oriented, operations-oriented, innovation, results-oriented, adaptability, integrity, and risk control; and management team for teamwork, innovation, results-oriented, adaptability, integrity, risk control, and people-oriented values. In terms of business outcomes shaped by corporate culture, analysts identify innovation and adaptability as affecting almost all aspects of business operations, ranging from market share and growth to corporate ESG practices, while other values such as customer-oriented, operation-oriented, or people-oriented have less impact on business outcomes. We further provide evidence that analysts' views of culture are distinct from values presented on corporate websites, and/or from the views of executives and employees.

Finally, we show that analysts' views of corporate culture are reflected in their stock recommendations and target price forecasts as well as impact price reactions to the release of their reports. We conclude that analysts' research on corporate culture offers new insights into its causes and effects, and that we are closer to establishing the culture-firm value link.

Our paper highlights the tremendous potential of generative AI in extracting cause-effect relations, and offers a roadmap for applying generative AI to finance and accounting research.

Appendix
Variable definitions

All continuous variables are winsorized at the 1st and 99th percentiles. All dollar values are in 2020 dollars.

Variable	Definition
Firm-year level	
Culture discussion	An indicator variable that takes the value of one if corporate culture is discussed in analyst reports in a year, and zero otherwise.
Number of segments	Number of culture-related segments in analyst reports in a year.
Number of reports	Number of analyst reports in which corporate culture is discussed in a year.
Number of values	Number of corporate cultural values discussed in analyst reports in a year.
Tone	The average tone of culture-related segments in analyst reports in a year. ChatGPT classifies each segment as negative (-1), neutral (0), or positive (1).
Total assets	Book value of total assets (in billions of dollars).
Firm size	Natural logarithm of total assets.
Firm age	Number of years since a firm first appears in Compustat.
Sales growth	One year sales changes divided by last year sales.
ROA	Operating income before interest and taxes divided by total assets.
Leverage	Book value of debt divided by total assets.
Tangibility	Net property, plant, and equipment divided by total assets.
ROA volatility	Standard deviation of annual ROA, calculated over last three years, multiplied by one hundred.
Large institutional ownership	The fraction of shares outstanding held by institutional investors with at least 5% of ownership of a firm. Missing values are assigned zero.
Board independence	The share of independent directors on a board.
CEO duality	An indicator variable that takes the value of one if the CEO is also Chairman of the Board, and zero otherwise.
Number of key people changes	Number of top executive and board member changes in a year. The data is from Capital IQ Key Developments database.
Number of M&As	Number of announcements related to mergers and acquisitions in a year. The data is from Capital IQ Key Developments database.
Corporate culture	The sum of five cultural values (innovation, integrity, quality, respect, and teamwork), extracted from earnings conference calls in a year, multiplied by 100. The data is from Li et al. (2021).
Employee culture rating	The average of employee ratings of culture & values (1-5) from all (i.e., current and former) employees in a year. The data is from Glassdoor.
Number of employee reviews	The number of all (i.e., current and former) employees providing culture & values ratings in a year. The data is from Glassdoor.
Adaptability	An indicator variable that takes the value of one if adaptability culture is discussed in analyst reports in a year, and zero otherwise. Other cultural value-specific indicators (e.g., customer-oriented) are defined analogously.
Firm-analyst-year level	
Culture discussion	An indicator variable that takes the value of one if corporate culture is discussed by an analyst in her reports in a firm-year, and zero otherwise.
Number of values	Number of corporate cultural values discussed by an analyst in her reports in a firm-year.
Tone	The average tone of culture-related segments in an analyst's reports in a firm-year.
Star analyst	An indicator variable that takes the value of one if an analyst is accredited to All-America research team status, and zero otherwise.
Female	An indicator variable that takes the value of one if an analyst is a female, and zero otherwise.

Forecast horizon	The average of forecast horizons (in terms of the number of years based on FPI in I/B/E/S) that an analyst employs when making forecasts in a firm-year.
General experience	The number of years for which an analyst makes at least one forecast of any firm.
Firm experience	The number of years for which an analyst makes at least one forecast of a given firm.
Number of industries followed	Number of two-digit SIC industries in which an analyst makes at least one forecast of any firm in that industry.
Number of firms followed	Number of firms for which an analyst makes at least one forecast.
Forecast frequency	Number of forecasts that an analyst makes of a given firm.
Broker size	Natural logarithm of number of analysts making at least one forecast at a given broker.
Report-level	
Recommendation	Stock recommendation in a report using a five-tier rating system where 2 represents "strong buy," 1 represents "buy," 0 represents "hold," -1 represents "underperform," and -2 represents "sell."
Recommendation down	An indicator variable that takes the value of one if stock recommendation in a report is revised downward compared to the last recommendation in I/B/E/S issued by the same analyst for the same firm, and zero if it is revised upward. To identify a revision, following Huang, Zang, and Zheng (2014, Figure 2), we first match a report to its corresponding I/B/E/S recommendation. A recommendation is considered valid during the period from the I/B/E/S announcement date until the I/B/E/S review date (i.e., when I/B/E/S confirms this recommendation is accurate). We classify a recommendation as revised if the report date is within the window from two days before to two days after the I/B/E/S announcement date, and a recommendation as reiterated if the report date is within the window from two days after the I/B/E/S announcement date to two days after the I/B/E/S review date.
Target price	Target price in a report divided by the stock price 50 days before the report date (in percentage points), following Huang, Zang, and Zheng (2014).
Target price down	An indicator variable that takes the value of one if target price in a report is revised downward compared to the last target price in I/B/E/S issued by the same analyst for the same firm, and zero if it is revised upward. To identify a revision, following Huang, Zang, and Zheng (2014, Figure 2), we first match a report to its corresponding I/B/E/S target price. A target price is considered valid during the period from the I/B/E/S announcement date until the I/B/E/S review date (i.e., when I/B/E/S confirms this estimate is accurate). We classify a target price as revised if the report date is within the window from two days before to two days after the I/B/E/S announcement date, and a recommendation as reiterated if the report date is within the window from two days after the I/B/E/S announcement date to two days after the I/B/E/S review date.
Tone	The average tone of culture-related segments in a report.
Non-culture tone	The average tone of non-culture-related segments in a report. FinBERT classifies each segment as negative (-1), neutral (0), or positive (1).
Adaptability tone	The average tone of adaptability-related segments in a report. ChatGPT classifies each segment as negative (-1), neutral (0), or positive (1). Other cultural value-specific (e.g., customer-oriented) tone variables are defined analogously.
Report length	Natural logarithm of number of segments in a report.
CAR[-1,+1]	Cumulative three-day abnormal return (in percentage points) centered around the report date (day 0) based on a market model in which the market portfolio is the CRSP value-weighted market index.
Earnings forecast revision	Earnings forecast in a report minus the last earnings forecast in I/B/E/S issued by the same analyst for the same firm, divided by the stock price 50 days before the report date.
Recommendation revision	Recommendation in a report minus the last recommendation in I/B/E/S issued by the same analyst for the same firm.

Target price revision	Target price in a report minus the last target price in I/B/E/S issued by the same analyst for the same firm, divided by the stock price 50 days before the report date.
Prior CAR	Cumulative ten-day abnormal return (in percentage points) ending two trading days before the report date based on a market model in which the market portfolio is the CRSP value-weighted market index.

References

Ahern, Kenneth R., Daniele Daminelli, and Cesare Fracassi, 2015, Lost in translation? The effect of cultural values on mergers around the world, *Journal of Financial Economics* 117, 165–189.

Antweiler, Werner, and Murray Z. Frank, 2004, Is all that talk just noise? The information content of internet stock message boards, *Journal of Finance* 59, 1259–1294.

Asquith, Paul, Michael B. Mikhail, and Andrea S. Au, 2005, Information content of equity analyst reports, *Journal of Financial Economics* 75, 245–282.

Au, Shiu-Yik, Ming Dong, and Andreanne Tremblay, 2021, Employee flexibility, exogenous risk, and firm value, *Journal of Financial and Quantitative Analysis* 56, 853–884.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* in press.

Bai, John Jianqiu, Nicole M. Boyson, Yi Cao, Miao Liu, and Chi Wan, 2023, Executives vs. chatbots: Unmasking insights through human-AI differences in earnings conference Q&A, Northeastern University working paper.

Bellstam, Gustaf, Sanjai Bhagat, and J. Anthony Cookson, 2021, A text-based analysis of corporate innovation, *Management Science* 67, 4004–4031.

Bertomeu, Jeremy, Yupeng Lin, Yibin Liu, and Zhenghui Ni, 2023, Capital market consequences of generative AI: Early evidence from the ban of ChatGPT in Italy, Washington University in St. Louis working paper.

Birru, Justin, Sinan Gokkaya, Xi Liu, and René M. Stulz, 2022, Are analyst short-term trade ideas valuable? *Journal of Finance* 77, 1829–1875.

Blanco, Eduardo, Nuria Castell, and Dan Moldovan, 2008, Causal relation extraction, *Proceedings of the 6th International Conference on Language Resources and Evaluation, LREC*.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan, 2003, Latent dirichlet allocation, *Journal of Machine Learning Research* 3, 993–1022.

Bradshaw, Mark T., Lawrence D. Brown, and Kelly Huang, 2013, Do sell-side analysts exhibit differential target price forecasting ability? *Review of Accounting Studies* 18, 930–955.

Brav, Alon, and Reuven Lehavy, 2003, An empirical analysis of analysts' target prices: Short-term informativeness and long-term dynamics, *Journal of Finance* 58, 1933–1967.

Briscoe-Tran, Hoa, 2022, Diversity and inclusion, operating flexibility, and financial policies, Ohio State University working paper.

Brown, Lawrence D., Andrew C. Call, Michael B. Clement, and Nathan Y. Sharp, 2015, Inside the "black box" of sell-side financial analysts, *Journal of Accounting Research* 53, 1–47.

Bybee, Leland, 2023, The ghost in the machine: Generating beliefs with large language models, ArXiv preprint arXiv:2305.02823.

Chi, Feng, Byoung-Hyoun Hwang, and Yaping Zheng, 2023, The use and usefulness of big data in finance: Evidence from financial analysts, Nanyang Business School Research Paper No. 22-01.

Choi, Freddy Y.Y., 2000, Advances in domain independent linear text segmentation, In *Proceedings of the 1st North American Chapter of the Association for Computational Linguistics Conference, NAACL 2000*, USA: Association for Computational Linguistics, 26–33.

Clement, Michael B., and Senyo Y. Tse, 2005, Financial analyst characteristics and herding behavior in forecasting, *Journal of Finance* 60, 307–341.

Derrien, François, and Ambrus Kecskés, 2013, The real effects of financial shocks: Evidence from exogenous changes in analyst coverage, *Journal of Finance* 68, 1407–1440.

Devlin, Jacob, Ming-wei Chang, Kenton Lee, and Kristina Toutanova, 2018, BERT: Pre-training of deep bi-directional transformers for language understanding, ArXiv preprint arXiv:1810.04805.

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang, 2023, Generative AI and firm values, University of California, Los Angeles working paper.

Etzioni, Oren, Michele Banko, Stephen Soderland, and Daniel S. Weld, 2008, Open information extraction from the web, *Communications of the ACM* 51, 68–74.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, 2019, Text as data, *Journal of Economic Literature* 57, 535–574.

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

Gorton, Gary B., Jillian Grennan, and Alexander K. Zentefis, 2022, Corporate culture, *Annual Review of Financial Economics* 14, 535–561.

Graham, John R., Jillian Grennan, Campbell R. Harvey, and Shivaram Rajgopal, 2022a, Corporate culture: Evidence from the field, *Journal of Financial Economics* 146, 552–593.

Graham, John R., Jillian A. Grennan, Campbell R. Harvey, and Shivaram Rajgopal, 2022b, Corporate culture: The interview evidence, *Journal of Applied Corporate Finance* 34, 22–41.

Grennan, Jillian, 2019, A corporate culture channel: How increased shareholder governance reduces firm value, Duke University working paper.

Grennan, Jillian, and Kai Li, 2023, Corporate culture: A review and directions for future research, In *Handbook of Financial Decision Making*, edited by Giles Hilary and David McLean, United Kingdom: Edward Elgar Publishing Ltd., 112–132.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2015, The value of corporate culture, *Journal of Financial Economics* 117, 60–76.

Heindorf, Stefan et al., 2020, CauseNet: Towards a causality graph extracted from the Web, *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, October 2020, 3023–3030.

Huang, Allen H., Reuven Lehavy, Amy Y. Zang, and Rong Zheng, 2018, Analyst information discovery and interpretation roles: A topic modeling approach, *Management Science* 64, 2833–2855.

Huang, Allen H., Amy Y. Zang, and Rong Zheng, 2014, Evidence on the information content of text in analyst reports, *Accounting Review* 89, 2151–2180.

Huang, Allen H., Hui Wang, and Yi Yang, 2023, FinBERT—A large language model approach to extracting information from financial text, *Contemporary Accounting Research* 40, 806–841.

Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang, 2023, ChatGPT and corporate policies, Georgia State University working paper.

Ji, Ziwei et al., 2023, Survey of hallucination in natural language generation, *ACM Computing Surveys* 55, 1–38.

Kecskés, Ambrus, Roni Michaely, and Kent L. Womack, 2017, Do earnings estimates add value to sell-side analysts' investment recommendations? *Management Science* 63, 1855–1871.

Kim, Alex G., Maximilian Muhn, and Valeri V. Nikolaev, 2023, Bloated disclosures: Can ChatGPT help investors process information? University of Chicago working paper.

Lewis, Patrick et al., 2020, Retrieval-augmented generation for knowledge-intensive NLP tasks, In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA: Curran Associates Inc., 9459–9474.

Li, Edward Xuejun, Zhiyuan Tu, and Dexin Zhou, 2023, The promise and peril of generative AI: Evidence from ChatGPT as sell-side analysts, City University of New York working paper.

Li, Feng, 2010, The information content of forward-looking statements in corporate filings—A Naïve Bayesian machine learning approach, *Journal of Accounting Research* 48, 1049–1102.

Li, Hongyi, and Eric Van den Steen, 2021, Birds of a feather … enforce social norms? Interactions among culture, norms, and strategy, *Strategy Science* 6, 166–189.

Li, Kai, Xing Liu, Feng Mai, and Tengfei Zhang, 2021, The role of corporate culture in bad times: Evidence from the COVID-19 pandemic, *Journal of Financial and Quantitative Analysis* 56, 2545–2583.

Li, Kai, Feng Mai, Rui Shen, and Xinyan Yan, 2021, Measuring corporate culture using machine learning, *Review of Financial Studies* 34, 3265–3315.

Li, Xina, Guoli Chen, and Rui Shen, 2024, Steering the intangible wheel: CEO succession and corporate cultural change, *Organization Science* forthcoming.

Liu, Nelson F. et al., 2023. Lost in the middle: How language models use long contexts, ArXiv preprint arXiv:2307.03172.

Liu, Xiaoding, 2016, Corruption culture and corporate misconduct, *Journal of Financial Economics* 122, 307–327.

Loh, Roger K., and René M. Stulz, 2011, When are analyst recommendation changes influential? *Review of Financial Studies* 24, 593–627.

Lopez-Lira, Alejandro, and Yuehua Tang, 2023, Can ChatGPT forecast stock price movements? Return predictability and large language models, University of Florida working paper.

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

Maynez, Joshua, Shashi Narayan, Bernd Bohnet, and Ryan McDonald, 2020, On faithfulness and factuality in abstractive summarization, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 1906–1919.

Mihăilă, Claudiu, Tomoko Ohta, Sampo Pyysalo, and Sophia Ananiadou, 2013, BioCause: Annotating and analysing causality in the biomedical domain, *BMC Bioinformatics* 14, 1–18.

Oh, Jong-Hoon et al., 2013, Why-question answering using intra-and inter-sentential causal relations, *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1733–1743.

Ouyang, Long et al., 2022, Training language models to follow instructions with human feedback, *Advances in Neural Information Processing Systems 35.*

Pacelli, Joseph, Tianshuo Shi, and Yuan Zou, 2022, Communicating corporate culture in labor markets: Evidence from job postings, Harvard Business School working paper.

Previts, Gary John, Robert J. Bricker, Thomas R. Robinson, and Stephen J. Young, 1994. A content analysis of sell-side financial analyst company reports, *Accounting Horizons* 8, 55–70.

Radinsky, Kira, Sagie Davidovich, and Shaul Markovitch, 2012, Learning causality for news events prediction, *Proceedings of the 21st International Conference on World Wide Web*, 909–918.

Reimers, Nils, and Iryna Gurevych, 2019, Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3982–3992.

Soltes, Eugene, 2014, Private interaction between firm management and sell-side analysts, *Journal of Accounting Research* 52, 245–272.

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

Twedt, Brady, and Lynn Rees, 2012, Reading between the lines: An empirical examination of qualitative attributes of financial analysts' reports, *Journal of Accounting and Public Policy* 31, 1–21.

Vaswani, Ashish et al., 2017, Attention is all you need, *Advances in Neural Information Processing Systems* 30.

Wei, Jason et al., 2022a, Chain-of-thought prompting elicits reasoning in large language models, *Advances in Neural Information Processing Systems* 35, 24824–24837.

Wei, Jason et al., 2022b, Emergent abilities of large language models, *Transactions on Machine Learning Research*.

Womack, Kent L., 1996. Do brokerage analysts' recommendations have investment value? *Journal of Finance* 51, 137–167.

Zingales, Luigi, 2015, The "cultural revolution" in finance, *Journal of Financial Economics* 117, 1–4.

Figure 1
Flowchart of our information extraction method built on generative AI

This figure presents a flowchart to implement our information extraction method built on generative AI models applied to millions of analyst reports over the period 2000–2020.

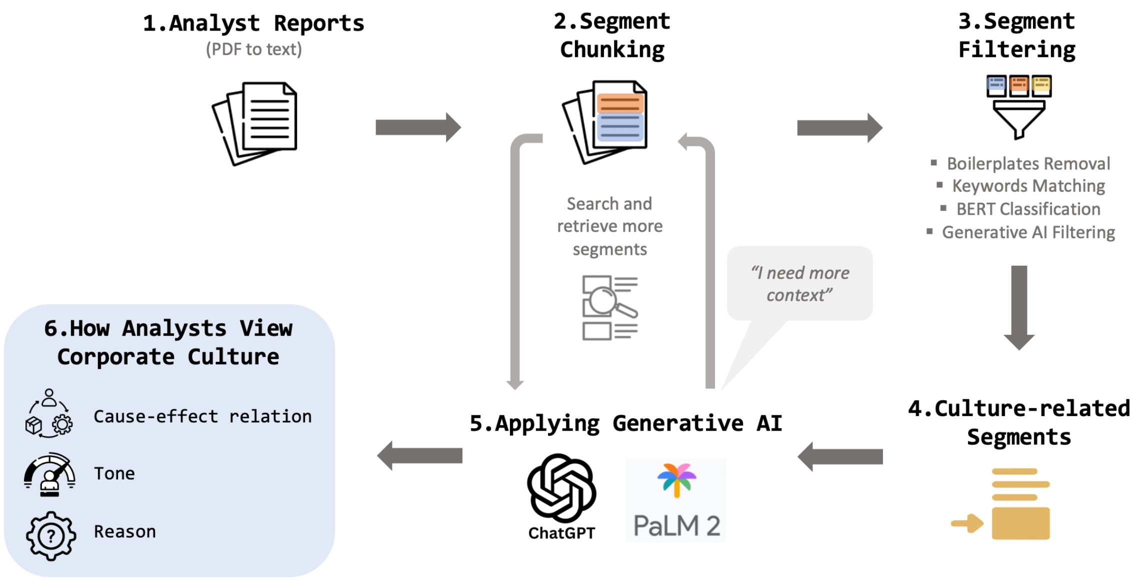

Figure 2
Intensity of analysts' culture-related discussions by year and report section

This heatmap depicts the intensity of analysts' culture-related discussions across years and report sections. Our sample comprises 2.4 million analyst reports over the period 2000–2020. The horizontal axis indicates report year, and the vertical axis indicates report section, binned into 20 equal sections from the start to the end of a report. The color gradient, ranging from light to dark, signifies the intensity of analysts' culture-related discussions. Intensity is computed as 100 × the number of culture-related segments in a section divided by the number of reports in a year.

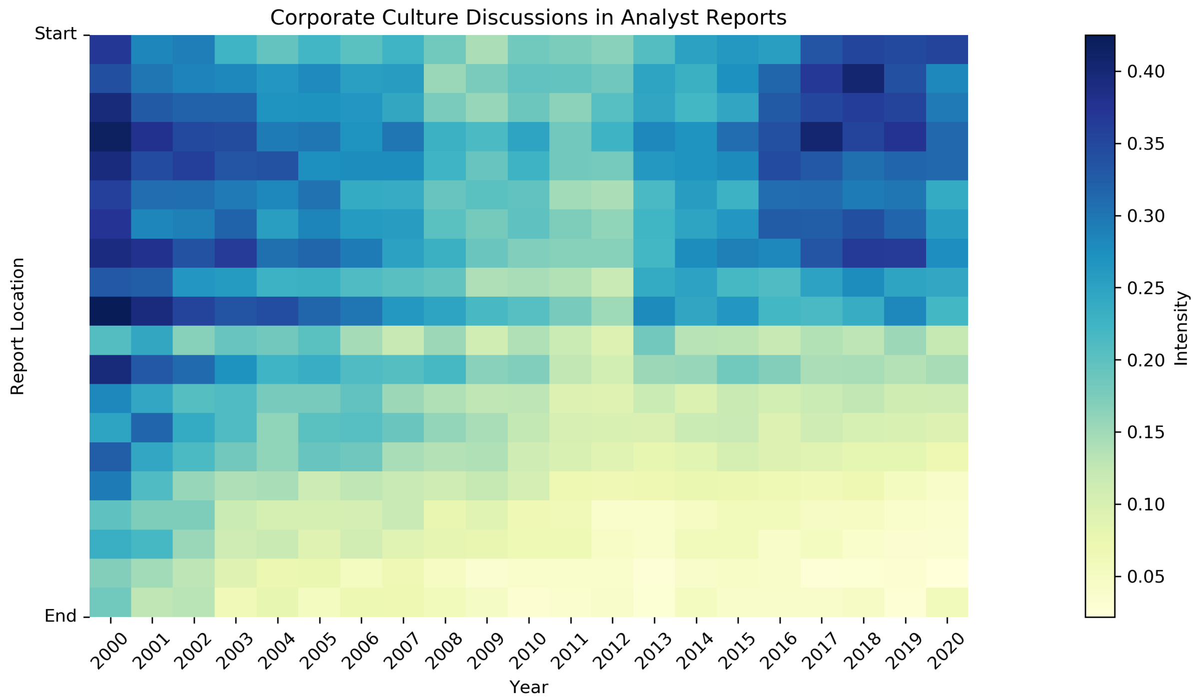

Figure 3
Cause-effect knowledge graph of corporate culture

This figure summarizes major cause-effect relations involving corporate culture. Our sample comprises 2.4 million analyst reports over the period 2000–2020. In the left column, we group the 17 causes into three groups: events, people, and systems. In the center column, we list nine cultural values. In the right column, we color-code the 16 business outcomes by tone. The height of each entity (cause, cultural value, or effect) corresponds to the number of relevant segments. The width of each link corresponds to the number of segments mentioning a particular cause or effect relation.

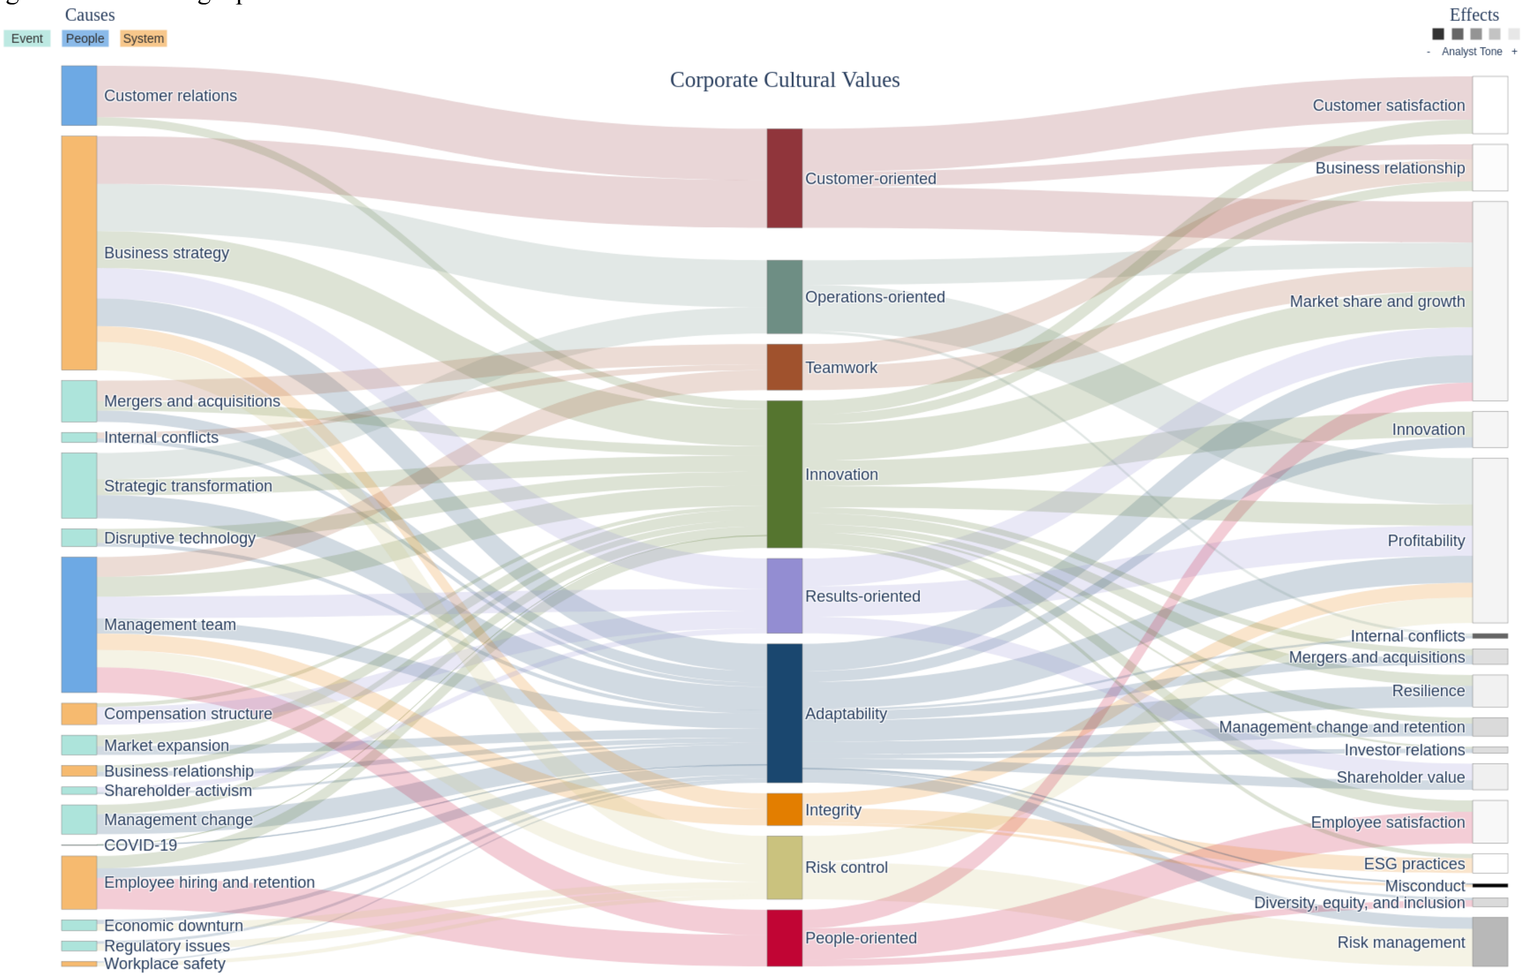

Table 1
Prompts for generative AI models

This table presents detailed instructions given to generative AI models to analyze and canonicalize information from corporate culture-related segments in analyst reports. Panel A shows the main prompt outlining the step-by-step process that generative AI models follow to extract information from each culture-related segment. It starts with identifying a cultural value (e.g., adaptability and teamwork), followed by extracting information on a value's influencing factor, its effect, and tone. The final step involves generating a cause-effect knowledge graph in the form of triples with their representation in a standard digital format called JSON (JavaScript Object Notation). Panel B shows the additional prompt for scenarios when more context is asked by generative AI models and is provided through retrieval augmented generation (RAG); under such scenarios, the underlined sections in Panel A are omitted. Panel C shows the prompt to canonicalize cultural values. Panel D shows the prompt to canonicalize causes and effects of a cultural value. Panel E shows the prompt to determine the direction of a cause or an effect. Panel F shows the prompt to canonicalize reasons for analysts to discuss culture. Text within *** is a placeholder that is adjusted dynamically.

Panel A: Chain-of-thought (CoT) prompt

```
As a sell-side equity analyst specializing in corporate culture and an expert in causal reasoning,
your task is to analyze a segment about corporate culture from a sell-side equity analyst research
report on a company. Your goal is to extract and interpret information about the company's corporate
culture, even if it is not explicitly stated. Your final goal is to extract cause-effect
relationships in a JSON format. Let's think step by step.

1. Identify the corporate cultural value being discussed. Summarize it in a short phrase starting
with an adjective. If the cultural value is not explicitly mentioned, infer it from the context.
Avoid using generic adjectives such as strong/weak/positive/negative culture. If uncertain, state "I
need more context" in the JSON value of the key.
2. Infer the main reason that corporate culture is discussed in the segment by the analyst. List the
main reason for discussing corporate culture, or mark as 'N/A' if not applicable.
3. If explicitly mentioned or can be inferred, summarize any events or factors that have shaped,
changed, or will change this corporate cultural value. List the most important causes of this
corporate cultural value, or mark as 'N/A' if not applicable.
4. If explicitly mentioned or can be inferred, summarize the past, present, or future outcomes or
impacts of this corporate cultural value on the company. List the most important outcomes of this
corporate cultural value, or mark as 'N/A' if not applicable.
5. Determine the tone of the analyst's discussion about the corporate cultural value. Options:
["positive", "negative", "neutral"].
6. Finally, based on your answers from previous steps, present a list of cause-effect graph triples
(those pertaining to the corporate cultural value) that you have extracted from the segment, or mark
as 'N/A' if not applicable.

If more context from the report is needed for the analysis, output "I need more context" in the JSON
value of the key.

Structure the output in a JSON format:

{
    "corporate cultural value": "corporate cultural value in a short phrase starting with an
adjective" or "I need more context",
    "reason for discussing corporate culture": "reason" or "N/A" or "I need more context",
    "causes of corporate culture": ["cause 1", ...] or "N/A" or "I need more context",
    "outcomes from corporate culture": ["outcome 1", ...] or "N/A" or "I need more context",
    "tone": "positive/negative/neutral",
    "cause-effect graph triples": [["entity_1", "relation", "entity_2"], ...] or "N/A",
}

When constructing "cause-effect graph triples", adhere to these criteria:
* "entity_1" or "entity_2" must be the identified cultural value.
* The other entity is either an outcome or a cause of corporate culture.
* "relation" must be a clear and simple verb phrase that conveys the direction of a cause-effect
relationship.
```

Panel B: Prompt to inform generative AI models that additional context is provided

```
######### Segment on corporate culture:
***input***
######### Additional contexts in the report (use only if needed):
***additional contexts***
######### Output:
```

Panel C: Prompt to canonicalize cultural values

```
Given a set of corporate cultural values as inputs, classify them into nine major categories using
the following definitions and examples as a guide.

* Adaptability: Cultures emphasizing cultural shifts, adaptations, transformations, or resistance to
change. For example, “cultural change,” “cultural transformation,” “culture reset,” and “insular
culture.”

* Customer-oriented: Cultures prioritizing sales or customer relations. For example, “sales
culture,” “customer-centric,” and “brand-focused.”

* Innovation: Cultures emphasizing innovation, growth, and entrepreneurship. For example,
“entrepreneurial,” “innovation,” “growth-oriented,” and “technology-driven.”

* Integrity: Cultures valuing ethical behavior, fair practices, and community relationships. For
example, “enlightened hospitality,” “ESG-focused,” “fair business practices,” and “blame culture.”

* Operations-oriented: Cultures underscoring efficiency, productivity, and cost control. For
example, “decentralized,” “efficiency-focused,” and “cost-conscious.”

* People-oriented: Cultures centering on employee well-being, diversity, inclusion, empowerment, and
talent growth. For example, “employee-centric,” “diverse and inclusive,” and “talent-focused.”

* Results-oriented: Cultures focusing on performance, competitiveness, and results. For example,
“competitive,” “performance-driven,” and “pay-for-performance.”

* Risk control: Cultures stressing risk management, credit, and financial prudence. For example,
“conservative credit culture,” “risk-averse,” and “culture of compliance.”

* Teamwork: Cultures highlighting collaboration, integration, and team-orientation. For example,
“collaborative,” “integration,” “cultural fit,” and “partnership-oriented.”

Thoroughly analyze each corporate cultural value and try your best to assign it to the most fitting
category, using "Miscellaneous" only when other categories do not apply.

* Miscellaneous: This category contains various unspecific and less-frequently occurring cultural
values, or those that do not fit easily into the above categories. For example, “positive culture”
and “distinct culture.”

Format your response in JSON. Make sure you process all of the inputs. Response format:

{"all results":
    [
        {
            "input_value_id": 1,
            "input_cultural_value": **example value**,
            "classification": **example classification**,
        },
        {
            "input_value_id": 2,
            "input_cultural_value": **example value**,
            "classification": **example classification**,
        },
        ... (other inputs)
    ]
}
```

Panel D: Prompt to canonicalize causes and effects of a corporate cultural value

```
Conduct entity canonicalization on selected phrases taken from sell-side equity analyst research
reports. These phrases are considered to be causes/consequences of a corporate cultural value. Your
objective is to classify each cause into one of {n_cats} predefined categories. If an entity does
not fit into these categories, the output should be "Other".

The available categories are:
----
{cats}
```

```
----

Format your response in JSON. Make sure you process all of the inputs. Response format and
examples:{"all results":
    [
        {
            "input_causes/consequence_id": 1,
            "input_causes/consequence": **example cause/consequence**,
            "classification": **example classification**,
        },
        {
            "input_causes/consequence_i": 2,
            "input_causes/consequence": **example cause/consequence**,
            "classification": **example classification**,
        },
        ... (other inputs)
    ]
}
```

Panel E: Prompt to canonicalize cause-effect directions

```
You are conducting cause-effect knowledge graph canonicalization, with a focus on relationships that
indicate a cause or an effect. Your assignment is as follows:

Establish the direction of the input relationship. You should map the direction into 'A -> B' or 'A
<- B', where the arrow indicates the direction of the cause-effect relationship. For clearly
unidirectional relationships such as 'affect', 'impact', 'influenced by', these should be marked as
'A -> B' or 'A <- B'. Only in rare cases where the direction of the relationship could be truly
bidirectional, mark the direction as 'A <-> B'.

Format your response in JSON. Make sure you process all of the input segments. Response format:

{
  "all results": [
    {
      "input_relationship_id": 1,
      "input_relationship": "A provides opportunity for B",
      "direction": "A -> B",
    },
    {
      "input_relationship_id": 2,
      "input_relationship": "A threatened by B",
      "direction": "A <- B",
    },
    {
      "input_relationship_id": 3,
      "input_relationship": "A influenced by B",
      "direction": "A <- B",
    },
    {
      "input_relationship_id": 4,
      "input_relationship": "A affects B",
      "direction": "A -> B",
    },
    ... (other inputs)
  ]
}
```

Panel F: Prompt to canonicalize reasons for discussing corporate culture

```
Analyze a set of reasons related to a sell-side equity analyst’s discussion of corporate culture in
their research reports on companies. Classify each input reason into one of the following five major
categories using the following definition provided for each category as a guide. If a reason does
not fit into any of the five categories, classify it as "Uncategorized".

* Financial performance, valuation, and competitive advantage: Encompass reasons focusing on the
impact of corporate culture on financial performance, competitive positioning, and market
differentiation. Analyze how corporate culture influences a company's financial health and
competitive edge.
* Mergers and acquisitions integration: Encompass reasons related to M&A activities, integration of
acquisitions, integration of mergers, and cultural impacts post-M&A. Assess how cultural integration
affects the potential synergies or challenges of these actions.
```

```
* Human capital management: Encompass reasons concerning talent attraction, retention, and employee
morale. Evaluate how corporate culture affects the company's ability to manage and develop its human
capital.
* Strategic alignment and execution: Encompass reasons focusing on strategic initiatives, management
changes, and organizational restructuring. Examine the alignment of corporate culture with the
company's strategy and its execution capabilities.
* Corporate governance and risk management: Encompass reasons related to corporate boards,
leadership dynamics, top management succession planning, and regulatory compliance. Assess how
corporate culture impacts corporate governance practices and risk management.

Format your response in JSON. Make sure you process all of the inputs. Response format:

{"all results":
    [
        {
            "input_reason_id": 1,
            "input_reason": "",
            "classification": "One of the 5 major categories",
        },
        {
            "input_reason_id": 2,
            "input_reason": "",
            "classification": "One of the 5 major categories",
        },
        ... (other inputs)
    ]
}
```

Table 2
Performance evaluation of different generative AI models

This table presents our performance evaluation of different generative AI models in terms of extracting cultural values, causes, effects, cause-effect, and tones. We compare the performance of GPT-4, GPT-3.5, and PaLM in terms of accuracy, precision, recall, and F1 against human annotations of a randomly chosen set of 200 culture-related segments. We group our fact-checking into four scenarios. True positive denotes a scenario in which generative AI extracts similar information about cultural values/causes/effects/cause-effect relations as we do. False positive denotes a scenario in which generative AI extracts different (false) information from we do. True negative denotes a scenario in which generative AI extracts no information, and neither do we. False negative denotes a scenario in which generative AI extracts false information while we do not find relevant information. We compute four performance metrics. Accuracy is defined as (#True Positive + #True Negative)/(#True Positive + #False Positive + #False Negative + #True Negative), and measures how accurate a model is at correctly classifying culture-related information out of the 200 segments. Precision is defined as (#True Positive)/(#True Positive + #False Positive) and measures how accurate a model is at identifying correct (positive) culture-related information out of all culture-related information that is predicted to be positive. Recall is defined as (#True Positive)/(#True Positive + #False Negative), and measures how accurate a model is at identifying correct (positive) culture-related information out of all identified culture-related information. F1 is the harmonic mean of Precision and Recall. Tone refers to the tone (negative, neutral, or positive) of culture-related information.

Performance metric	AI Model	Culture value	Cause	Effect	Cause-effect	Tone
Accuracy	GPT-4	94.5%	91.5%	94.0%	80.5%	93.0%
	GPT-3.5	91.0%	86.5%	88.5%	71.0%	91.5%
	PaLM	81.5%	80.0%	85.5%	67.0%	88.5%
Precision	GPT-4	94.5%	93.0%	94.9%	82.1%	93.0%
	GPT-3.5	91.0%	90.4%	92.6%	75.6%	91.5%
	PaLM	89.5%	88.3%	90.9%	74.5%	93.2%
Recall	GPT-4	100.0%	97.7%	98.9%	96.2%	100.0%
	GPT-3.5	100.0%	94.1%	95.1%	88.2%	100.0%
	PaLM	90.0%	87.3%	93.4%	81.8%	94.7%
F1	GPT-4	97.2%	95.3%	96.9%	88.6%	96.4%
	GPT-3.5	95.3%	92.2%	93.8%	81.4%	95.6%
	PaLM	89.8%	87.8%	92.1%	78.0%	93.9%

Table 3
Frequency distributions of culture-related discussions in analyst reports

This table presents frequency distributions of culture-related discussions (cultural values, causes, and effects) in analyst reports over the period 2000‒2020. Panel A lists the frequency distribution of cultural values in 32,186 reports with culture-related segments. Panel B lists the frequency distribution of causes of culture in 27,064 reports with culture-related segments. Panel C lists the frequency distribution of effects of culture in 31,643 reports with culture-related segments. One report could be associated with multiple cultural values, causes, or effects.

Panel A: Frequency distribution of cultural values

Cultural value	Count	Percentage
Innovation	15,601	16.9%
Adaptability	14,711	15.9%
Miscellaneous	14,690	15.9%
Customer-oriented	10,634	11.5%
Results-oriented	7,923	8.6%
Operations-oriented	7,800	8.4%
Risk control	6,788	7.3%
People-oriented	6,033	6.5%
Teamwork	4,873	5.3%
Integrity	3,419	3.7%
Total	92,472	100.0%

Panel B: Frequency distribution of causes of culture

Cause	Count	Percentage
Business strategy	24,874	18.4%
Management team	16,761	12.4%
Miscellaneous	15,255	11.3%
Strategic transformation	11,960	8.8%
Employee hiring and retention	10,519	7.8%
Mergers and acquisitions	9,314	6.9%
Management change	9,258	6.8%
Customer relations	7,111	5.3%
Market expansion	6,256	4.6%
Business relationship	5,296	3.9%
Compensation structure	4,818	3.6%
Disruptive technology	3,997	3.0%
Internal conflicts	2,489	1.8%
Economic downturn	2,168	1.6%
Shareholder activism	1,911	1.4%
Regulatory issues	1,720	1.3%
Workplace safety	1,299	1.0%
COVID-19	178	0.1%
Total	135,184	100.0%

Panel C: Frequency distribution of effects of culture

Effect	Count	Percentage
Market share and growth	26,172	16.2%
Profitability	21,934	13.6%
Employee satisfaction	13,519	8.4%
Customer satisfaction	12,941	8.0%
Resilience	12,815	7.9%
Innovation	10,806	6.7%
Risk management	9,779	6.1%
Business relationship	8,946	5.5%
Shareholder value	8,623	5.3%
Management change and retention	8,293	5.1%
Miscellaneous	8,250	5.1%
Mergers and acquisitions	7,450	4.6%
ESG practices	4,289	2.7%
Investor relations	3,062	1.9%
Diversity, equity, and inclusion	2,110	1.3%
Internal conflicts	1,581	1.0%
Misconduct	728	0.5%
Total	161,298	100.0%

Table 4
Sample formation

This table reports the impact of various data filters on sample formation. Our sample starts from Thomson One's Investext over the period 2002-2020.

Panel A: Sample formation of the firm-year sample

	# firm-year obs.	# firm-year obs. removed	# unique firms
Firm-years with analyst reports in Thomson One's Investext.	34,414		2,889
Remove firm-years with missing financial information from Compustat.	30,978	3,436	2,728
Remove firm-years with missing information from ExecuComp.	**28,860**	2,118	**2,569**
Remove firm-years without corporate culture data from Li et al. (2021).	24,608	4,252	2,329
Remove firm-years without Glassdoor data.	9,645	14,963	1,143

Panel B: Sample formation of the firm-analyst-year sample

	# firm-analyst-year obs.	# firm-analyst-year obs. removed	# unique firms
Firm-analyst-years with analyst reports in Thomson One's Investext.	243,856		2,907
Remove firm-analyst-years not in our firm-year sample.	147,180	96,676	2,490
Remove firm-analyst-years with missing analyst characteristics from I/B/E/S.	**142,868**	4,312	**2,484**

Table 5
Summary statistics

This table presents the summary statistics for samples used in different regression analyses. Our firm-year sample consists of 28,860 firm-year observations, representing 2,569 unique firms over the period 2002-2020. Our firm-analyst-year sample consists of 142,868 firm-analyst-year observations, representing 2,484 unique firms followed by 4,043 analysts. Definitions of the variables are provided in the Appendix.

	Mean	25th Percentile	Median	75th Percentile	SD
Firm-year sample					
Culture discussion	0.425	0.000	0.000	1.000	0.494
Number of values	6.736	2.000	10.000	10.000	3.953
Tone	0.676	0.500	1.000	1.000	0.509
Adaptability	0.151	0.000	0.000	0.000	0.358
Customer-oriented	0.093	0.000	0.000	0.000	0.291
Innovation	0.155	0.000	0.000	0.000	0.362
Integrity	0.045	0.000	0.000	0.000	0.208
Operations-oriented	0.090	0.000	0.000	0.000	0.286
People-oriented	0.072	0.000	0.000	0.000	0.258
Results-oriented	0.096	0.000	0.000	0.000	0.295
Risk control	0.071	0.000	0.000	0.000	0.257
Teamwork	0.062	0.000	0.000	0.000	0.241
Miscellaneous	0.151	0.000	0.000	0.000	0.358
Total assets	14.204	0.892	2.747	9.320	38.370
Firm age	26.566	13.000	21.000	39.000	17.360
Sales growth	0.097	-0.009	0.070	0.168	0.225
ROA	0.037	0.011	0.042	0.081	0.096
Leverage	0.234	0.065	0.211	0.352	0.195
Tangibility	0.237	0.051	0.149	0.357	0.238
ROA volatility	0.040	0.007	0.018	0.042	0.063
Large institutional ownership	0.369	0.110	0.349	0.579	0.294
Board independence	0.714	0.615	0.688	0.857	0.136
CEO duality	0.493	0.000	0.000	1.000	0.500
Number of key people changes	3.447	1.000	2.000	5.000	3.582
Number of M&As	0.713	0.000	0.000	1.000	1.199
Corporate culture	13.770	10.126	12.844	16.518	4.933
Employee culture rating	2.210	0.000	2.700	3.250	1.441
Number of employee reviews	88.950	4.000	18.000	68.000	216.707
Firm-analyst-year sample					
Culture discussion	0.121	0.000	0.000	0.000	0.326
Number of segments	1.900	1.000	1.000	2.000	1.327
Number of reports	1.569	1.000	1.000	2.000	0.956
Number of values	1.495	1.000	1.000	2.000	0.943
Tone	0.667	0.500	1.000	1.000	0.582
Star analyst	0.072	0.000	0.000	0.000	0.259
Female	0.110	0.000	0.000	0.000	0.313
Forecast horizon	1.237	1.035	1.152	1.334	0.316
General experience	11.030	5.000	10.000	15.000	6.973
Firm experience	5.342	2.000	4.000	7.000	4.387
Number of industries followed	4.482	2.000	4.000	6.000	2.596

Number of firms followed	19.009	14.000	18.000	23.000	8.168
Forecast frequency	4.737	3.000	4.000	6.000	2.475
Number of analysts at a broker	64.314	26.000	54.000	103.000	44.235

Table 6
Analysts discussing corporate culture in reports and executives discussing corporate culture in earnings calls

This table examines whether analysts' discussions of corporate culture are mainly driven by what they hear from earnings calls. The sample comprises 12,267 firm-year observations whose analysts discuss corporate culture in their reports over the period 2002–2020. Panel A presents two-sample t-test and Wilcoxon rank-sum test between firm-year observations with earnings conference calls and firm-year observations without. Panel B presents two-sample t-test and Wilcoxon rank-sum test between firm-year observations with high *Corporate culture* and firm-year observations with low *Corporate culture*. *Corporate culture* is the sum of five cultural values extracted from earnings conference calls in a year. The data is from Li et al. (2021). A firm is classified as having high (low) *Corporate culture* if its value is in the top (bottom) quartile in an industry-year. Industry is based on Fama-French 12-industry classifications. Definitions of the variables are provided in the Appendix. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Firm-year observations with and without earnings calls

	Subsample with earnings calls		Subsample without earnings calls		Two-sample differences	
	Mean	Median	Mean	Median	t-test	Wilcoxon test
Number of segments	4.406	2.000	4.842	2.000	-0.436**	0.000
Number of reports	3.190	2.000	3.073	2.000	0.117	0.000***
Number of values	2.313	2.000	2.369	2.000	-0.056	0.000
Tone	0.675	1.000	0.682	1.000	-0.007	-0.000
Adaptability	0.360	0.000	0.325	0.000	0.035***	0.000***
Customer-oriented	0.214	0.000	0.258	0.000	-0.044***	-0.000***
Innovation	0.372	0.000	0.310	0.000	0.062***	0.000***
Integrity	0.107	0.000	0.103	0.000	0.004	0.000
Operations-oriented	0.211	0.000	0.219	0.000	-0.008	-0.000
People-oriented	0.169	0.000	0.166	0.000	0.003	0.000
Results-oriented	0.222	0.000	0.256	0.000	-0.034***	-0.000***
Risk control	0.160	0.000	0.215	0.000	-0.055***	-0.000***
Teamwork	0.144	0.000	0.154	0.000	-0.010	-0.000
Miscellaneous	0.355	0.000	0.363	0.000	-0.008	-0.000
Total assets	23.248	5.206	24.393	4.408	-1.145	0.798**
Firm age	28.442	23.000	25.015	21.000	3.427***	2.000***
Sales growth	0.097	0.071	0.101	0.073	-0.004	0.002
ROA	0.051	0.050	0.045	0.041	0.006***	0.009***
Leverage	0.233	0.211	0.209	0.187	0.024***	0.024***
Tangibility	0.216	0.132	0.215	0.133	0.001	-0.001
ROA volatility	0.030	0.015	0.028	0.012	0.002**	0.003***
Large institutional ownership	0.367	0.353	0.234	0.160	0.133***	0.193***
Board independence	0.708	0.688	0.660	0.643	0.048***	0.045***
CEO duality	0.504	1.000	0.593	1.000	-0.089**	-0.000***

Panel B: Firm-year observations with high and low corporate culture

	Subsample with high corporate culture		Subsample with low corporate culture		Two-sample differences	
	Mean	Median	Mean	Median	t-test	Wilcoxon test
Number of segments	5.228	3.000	3.794	2.000	1.434***	1.000***
Number of reports	3.662	2.000	2.767	2.000	0.895***	0.000***
Number of values	2.540	2.000	2.084	1.000	0.456***	1.000***
Tone	0.651	1.000	0.685	1.000	-0.034**	-0.000***
Adaptability	0.406	0.000	0.321	0.000	0.085***	0.000***
Customer-oriented	0.249	0.000	0.180	0.000	0.069***	0.000***
Innovation	0.438	0.000	0.297	0.000	0.141***	0.000***
Integrity	0.120	0.000	0.102	0.000	0.018**	0.000**
Operations-oriented	0.209	0.000	0.218	0.000	-0.009	0.000
People-oriented	0.191	0.000	0.133	0.000	0.058***	0.000***
Results-oriented	0.238	0.000	0.224	0.000	0.014	0.000
Risk control	0.141	0.000	0.178	0.000	-0.037***	-0.000***
Teamwork	0.161	0.000	0.121	0.000	0.040***	0.000***
Miscellaneous	0.386	0.000	0.311	0.000	0.075***	0.000***
Total assets	22.370	3.407	21.940	6.626	0.43	-3.219***
Firm age	24.540	19.000	31.717	28.000	-7.177***	-9.000***
Sales growth	0.108	0.077	0.084	0.067	0.024***	0.010***
ROA	0.050	0.052	0.048	0.047	0.002	0.005**
Leverage	0.211	0.177	0.249	0.228	-0.038***	-0.051***
Tangibility	0.193	0.113	0.234	0.157	-0.041***	-0.044***
ROA volatility	0.036	0.019	0.026	0.013	0.010***	0.006***
Large institutional ownership	0.364	0.354	0.359	0.348	0.005	0.006
Board independence	0.705	0.688	0.714	0.688	-0.009**	-0.000***
CEO duality	0.472	0.000	0.536	1.000	-0.064***	-1.000***

Table 7
Determinants of analysts' discussing corporate culture in reports

This table examines the determinants of analysts' discussing corporate culture in their reports. Panel A examines the relations between firm characteristics and analysts' discussing culture at the firm-year level. Our firm-year sample consists of 28,860 firm-year observations, representing 2,569 unique firms over the period 2002-2020. The dependent variable, *Culture discussion*, is an indicator variable that takes the value of one if corporate culture is discussed in analyst reports in a year, and zero otherwise. Panel B examines the relations between firm characteristics and the scope and tone of analysts' discussing culture in reports. *Number of values* is the number of corporate cultural values discussed in analyst reports in a year. *Tone* is the average tone of culture-related segments in analyst reports in a year. Panel C examines the relations between analyst characteristics and them discussing culture at the firm-analyst-year level. Our firm-analyst-year sample consists of 142,868 firm-analyst-year observations, representing 2,484 unique firms followed by 4,043 analysts. Industry fixed effects are based on Fama-French 12-industry classifications. Definitions of the variables are provided in the Appendix. Heteroskedasticity-consistent standard errors (in parentheses) are clustered at the firm level. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Firm characteristics and analysts' discussing corporate culture in reports

	Culture discussion			
Variable	(1)	(2)	(3)	(4)
Firm size	0.100***	0.084***	0.086***	0.046***
	(0.004)	(0.004)	(0.004)	(0.007)
Ln(Firm age + 1)	-0.026***	-0.029***	-0.018**	-0.018
	(0.008)	(0.008)	(0.009)	(0.013)
Sales growth	0.031**	0.033**	0.028*	0.101***
	(0.014)	(0.014)	(0.015)	(0.027)
ROA	0.291***	0.329***	0.307***	0.144*
	(0.046)	(0.045)	(0.047)	(0.079)
Leverage	-0.163***	-0.156***	-0.122***	-0.032
	(0.027)	(0.026)	(0.027)	(0.041)
Tangibility	-0.092***	-0.078***	-0.057*	-0.067
	(0.030)	(0.029)	(0.030)	(0.042)
ROA volatility	-0.134**	-0.199***	-0.276***	-0.437***
	(0.063)	(0.062)	(0.066)	(0.116)
Large institution ownership	-0.074***	-0.071***	-0.070***	-0.099***
	(0.016)	(0.015)	(0.016)	(0.025)
Board independence	-0.195***	-0.227***	-0.239***	-0.402***
	(0.039)	(0.038)	(0.040)	(0.067)
CEO duality	0.004	0.011	0.011	0.009
	(0.009)	(0.009)	(0.009)	(0.014)
Ln(Number of key people changes + 1)		0.068***	0.055***	0.042***
		(0.005)	(0.006)	(0.009)
Ln(Number of M&As + 1)		0.028***	0.021***	0.010
		(0.007)	(0.008)	(0.012)
Corporate culture			0.014***	0.010***
			(0.001)	(0.002)
Employee culture rating				0.005
				(0.007)
Ln(Number of employee reviews + 1)				0.047***
				(0.006)
Constant	YES	YES	YES	YES
Industry FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Adjusted R^2	0.146	0.157	0.172	0.174
No. of observations	28,860	28,860	24,608	9,645

Panel B: Firm characteristics and the scope and tone of analysts' discussing corporate culture in reports

	Number of values				Tone			
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Firm size	0.311***	0.242***	0.244***	0.158***	-0.019***	-0.014***	-0.015***	-0.025***
	(0.026)	(0.025)	(0.023)	(0.040)	(0.005)	(0.005)	(0.005)	(0.009)
Ln(Firm age + 1)	-0.001	-0.009	0.030	0.036	-0.005	-0.004	-0.005	-0.009
	(0.046)	(0.044)	(0.044)	(0.064)	(0.010)	(0.010)	(0.011)	(0.015)
Sales growth	0.141*	0.175**	0.181**	0.453***	0.053*	0.049*	0.044	0.095**
	(0.084)	(0.083)	(0.088)	(0.148)	(0.029)	(0.029)	(0.031)	(0.046)
ROA	1.129***	1.283***	0.858***	-0.448	0.254***	0.241***	0.220**	0.205
	(0.312)	(0.302)	(0.279)	(0.480)	(0.085)	(0.086)	(0.090)	(0.129)
Leverage	-0.354***	-0.346***	-0.242*	-0.070	-0.012	-0.014	-0.014	0.036
	(0.130)	(0.125)	(0.129)	(0.187)	(0.033)	(0.033)	(0.035)	(0.049)
Tangibility	0.069	0.140	0.176	-0.091	0.081**	0.076**	0.070**	0.099**
	(0.165)	(0.159)	(0.158)	(0.203)	(0.033)	(0.033)	(0.035)	(0.047)
ROA volatility	-0.450	-0.858**	-1.245***	-1.541***	-0.422***	-0.387***	-0.268*	-0.031
	(0.358)	(0.350)	(0.339)	(0.586)	(0.147)	(0.146)	(0.146)	(0.205)
Large institution ownership	-0.320***	-0.313***	-0.285***	-0.270**	0.031	0.031	0.021	0.018
	(0.084)	(0.082)	(0.082)	(0.128)	(0.023)	(0.023)	(0.024)	(0.039)
Board independence	-0.118	-0.257	-0.175	-0.461	-0.025	-0.013	-0.001	-0.064
	(0.224)	(0.218)	(0.214)	(0.329)	(0.051)	(0.050)	(0.054)	(0.079)
CEO duality	-0.008	0.024	0.037	-0.007	0.016	0.014	0.020	-0.004
	(0.052)	(0.051)	(0.050)	(0.073)	(0.012)	(0.012)	(0.013)	(0.020)
Ln(Number of key people changes + 1)		0.274***	0.214***	0.187***		-0.024***	-0.017**	-0.013
		(0.030)	(0.031)	(0.046)		(0.007)	(0.008)	(0.012)
Ln(Number of M&As + 1)		0.165***	0.127***	0.128*		-0.009	-0.005	-0.004
		(0.045)	(0.046)	(0.070)		(0.010)	(0.011)	(0.016)
Corporate culture			0.050***	0.045***			-0.003**	-0.003*
			(0.006)	(0.009)			(0.001)	(0.002)
Employee culture rating				-0.001				-0.005
				(0.036)				(0.012)
Ln(Number of employee reviews + 1)				0.140***				0.016*
				(0.034)				(0.008)
Constant	YES	YES	YES	YES	YES	YES	YES	YES
Industry FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES	YES

Adjusted R^2	0.133	0.149	0.157	0.167	0.020	0.021	0.018	0.016
No. of observations	12,267	12,267	10,699	4,821	12,267	12,267	10,699	4,821

Panel C: Analyst characteristics and their discussing corporate culture in reports

	Culture discussion	Number of values	Tone	Culture discussion	Number of values	Tone
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Star analyst	-0.001	0.010	-0.035	0.019***	0.138***	-0.061**
	(0.005)	(0.049)	(0.024)	(0.005)	(0.051)	(0.026)
Female	0.010**	0.065	0.028	0.010**	0.069	0.020
	(0.004)	(0.046)	(0.018)	(0.004)	(0.045)	(0.019)
Forecast horizon	0.003	0.163***	-0.025	0.008*	0.135**	-0.002
	(0.004)	(0.052)	(0.024)	(0.004)	(0.066)	(0.031)
General experience	0.001***	0.006**	-0.000	0.001***	0.004*	-0.000
	(0.000)	(0.002)	(0.001)	(0.000)	(0.002)	(0.001)
Firm experience	0.002***	0.001	-0.001	0.002***	0.002	-0.000
	(0.000)	(0.003)	(0.002)	(0.000)	(0.003)	(0.002)
Number of industries followed	0.001*	0.020**	0.000	-0.001	0.014*	0.001
	(0.001)	(0.008)	(0.004)	(0.001)	(0.008)	(0.004)
Number of firms followed	-0.000**	-0.002	0.000	-0.000	-0.001	0.000
	(0.000)	(0.002)	(0.001)	(0.000)	(0.002)	(0.001)
Forecast frequency	-0.001	-0.005	-0.000	0.001**	0.003	-0.002
	(0.001)	(0.007)	(0.004)	(0.001)	(0.007)	(0.004)
Broker size	0.015***	0.071***	0.017*	0.013***	0.000	0.051
	(0.001)	(0.017)	(0.010)	(0.004)	(0.054)	(0.036)
Constant	YES	YES	YES	YES	YES	YES
Firm × Year FE	YES	YES	YES	YES	YES	YES
Broker FE	NO	NO	NO	YES	YES	YES
Adjusted R^2	0.114	0.050	0.144	0.136	0.083	0.162
No. of observations	142,868	17,309	17,309	142,868	17,309	17,309

Table 8
Analysts' views of corporate culture and their research output

This table examines the relationships between analysts' views of corporate culture and their stock recommendations and target prices at the report level. Panel A presents the summary statistics for the key variables. Panel B examines the relationships between analysts' tones in culture-related segments and their stock recommendations and target prices. Panel C examines the relationships between analysts' cultural value-specific tones in culture-related segments and their stock recommendations and target prices. The dependent variable in column (1), *Recommendation*, is a report's stock recommendation using a five-tier rating system where 2 represents "strong buy," 1 represents "buy," 0 represents "hold," -1 represents "underperform," and -2 represents "sell." The dependent variable in column (2), *Target price*, is a report's target price divided by the stock price 50 days before the report date (in percentage points). Panel D repeats the analysis in Panel B limiting to a subsample of reports with recommendation/target price revisions. Industry fixed effects are based on Fama-French 12-industry classifications. Definitions of the variables are provided in the Appendix. Heteroskedasticity-consistent standard errors (in parentheses) are double-clustered at the firm and analyst levels. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Summary statistics for the key variables

	Mean	25th Percentile	Median	75th Percentile	SD
Stock recommendation sample					
Recommendation	0.650	0.000	1.000	1.000	0.854
Recommendation down	0.413	0.000	0.000	1.000	0.492
Tone	0.016	0.000	0.000	0.000	0.142
Target price sample					
Target price (%)	118.141	104.454	116.908	129.921	23.910
Target price down	0.361	0.000	0.000	1.000	0.480

Panel B: The tone in culture-related segments, stock recommendations, and target prices

	Recommendation	Target price
Variable	(1)	(2)
Tone	0.109***	1.653***
	(0.009)	(0.198)
Non-culture tone	0.312***	6.882***
	(0.007)	(0.174)
Report length	0.014***	0.011
	(0.002)	(0.042)
Star analyst	0.011	0.109
	(0.015)	(0.403)
Female	0.000	0.000
	(0.000)	(0.000)
Forecast horizon	-0.010	0.052
	(0.022)	(0.494)
General experience	-0.001	-0.072
	(0.004)	(0.118)
Firm experience	0.006***	0.078***
	(0.001)	(0.029)
Number of industries followed	0.002	0.027
	(0.003)	(0.080)
Number of firms followed	-0.001*	-0.004
	(0.001)	(0.019)
Forecast frequency	0.006***	0.091**
	(0.001)	(0.039)
Broker size	-0.051***	-0.651**
	(0.012)	(0.259)

Constant	YES	YES
Firm × Year FE	YES	YES
Analyst FE	YES	YES
Adjusted R^2	0.450	0.399
No. of observations	1,040,758	1,033,669

Panel C: The tone in cultural value-related segments, stock recommendations, and target prices

	Recommendation	Target price
Variable	(1)	(2)
Adaptability tone	0.129***	1.824***
	(0.017)	(0.408)
Customer-oriented tone	0.106***	1.344***
	(0.021)	(0.454)
Innovation tone	0.103***	1.803***
	(0.018)	(0.487)
Integrity tone	0.122***	2.385***
	(0.033)	(0.755)
Operations-oriented tone	0.075***	1.073**
	(0.024)	(0.530)
People-oriented tone	0.131***	1.987**
	(0.049)	(0.776)
Results-oriented tone	0.098***	1.678**
	(0.025)	(0.680)
Risk control tone	0.111***	1.250**
	(0.026)	(0.566)
Teamwork tone	0.097***	3.071***
	(0.031)	(0.991)
Miscellaneous tone	0.116***	1.197**
	(0.018)	(0.477)
Non-culture tone	0.312***	6.882***
	(0.007)	(0.174)
Report length	0.014***	0.011
	(0.002)	(0.042)
Other analyst controls	YES	YES
Constant	YES	YES
Firm × Year FE	YES	YES
Analyst FE	YES	YES
Adjusted R^2	0.450	0.399
No. of observations	1,040,758	1,033,669

Panel D: The tone in culture-related segments, stock recommendation revisions, and target price revisions

	Recommendation down	Target price down
Variable	(1)	(2)
Tone	-0.074***	-0.041***
	(0.013)	(0.005)
Non-culture tone	-0.387***	-0.325***
	(0.007)	(0.004)
Report length	-0.025***	0.005***
	(0.003)	(0.001)
Other analyst controls	YES	YES
Constant	YES	YES
Firm × Year FE	YES	YES
Analyst FE	YES	YES
Adjusted R^2	0.261	0.361
No. of observations	60,441	339,875

Table 9
Information content of analysts' views of corporate culture

This table examines the information content of analysts' views of corporate culture at the report level. The sample comprises 47,336 reports that contain an earnings forecast revision and are not issued at the same time as other reports on the same firm or any other major corporate announcements. Panel A presents the summary statistics for the key variables. Panel B presents the regression results. The dependent variable, *CAR[-1,+1]*, is the cumulative abnormal return (in percentage points) centered around the report date (day 0) based on a market model. Industry fixed effects are based on Fama-French 12-industry classifications. Definitions of the variables are provided in the Appendix. Heteroskedasticity-consistent standard errors (in parentheses) are double-clustered at the firm and analyst levels. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Summary statistics for the key variables

	Mean	25^{th} Percentile	Median	75^{th} Percentile	SD
CAR[-1,+1] (%)	-0.024	-1.932	-0.047	1.856	4.489
Tone	0.019	0.000	0.000	0.000	0.156
Non-culture tone	0.174	-0.111	0.200	0.500	0.476

Panel B: Price reactions to analyst reports

	CAR[-1,+1]	CAR[-1,+1]
Variable	(1)	(2)
Tone	0.306***	
	(0.118)	
Adaptability tone		0.194
		(0.358)
Customer-oriented tone		0.393
		(0.434)
Innovation tone		0.229
		(0.247)
Integrity tone		0.757
		(0.805)
Operations-oriented tone		1.122***
		(0.413)
People-oriented tone		0.177
		(0.492)
Results-oriented tone		0.287
		(0.302)
Risk control tone		0.433
		(0.329)
Teamwork tone		-0.076
		(0.538)
Miscellaneous tone		0.026
		(0.348)
Non-culture tone	1.136***	1.136***
	(0.055)	(0.055)
Report length	0.065***	0.065***
	(0.025)	(0.025)
Earnings forecast revision	29.463***	29.463***
	(0.031)	(0.031)
Recommendation revision	0.264***	0.264***
	(0.019)	(0.019)
Target price revision	0.010***	0.010***
	(0.000)	(0.000)
Prior CAR	-0.020***	-0.020***
	(0.004)	(0.004)

Other analyst/firm controls	YES	YES
Constant	YES	YES
Industry FE	YES	YES
Year FE	YES	YES
Adjusted R^2	0.044	0.044
No. of observations	47,336	47,336

Internet Appendix for

Dissecting Corporate Culture Using Generative AI – Insights from Analyst Reports

A. Technical Appendix

In this appendix, we describe technical details of how we preprocess analyst reports, remove boilerplate segments, identify culture-related segments via word sense disambiguation and a machine learning model, and implement canonicalization of cultural values and causes/effects. Figure 1 provides a flowchart of our information extraction method built on generative AI models.

1. Converting Reports from PDF to Text

We download 2,434,782 reports over the period 2000–2020 from Thomson One's Investext database. The reports are in PDF format. We use GROBID (https://github.com/kermitt2/grobid), an open-source software, to extract structured information from PDF documents and transform this information into XML documents. The XML documents are then stripped of information identified as tables, annexes, notes, and author information; the main content is converted to plain text. We further split text into sentences using OpenNLP's sentence segment module, a built-in function in GROBID.

2. Segment Chunking

An inherent challenge resulting from the conversion process described above is the loss of paragraph structure in reports. Moreover, even if the structure could have been maintained, differentiating between headers, bullet points, and coherent paragraphs in a report is not straightforward.

To address these issues, we employ the C99 algorithm, a common text segmentation technique developed by Choi (2000). The C99 algorithm is a domain-independent, unsupervised method for linear text segmentation. Its defining principle is that topics in a text document are coherent, and topic shifts can be identified by a sharp decline in coherence. The algorithm quantifies coherence through pairwise similarity of sentences. It builds a matrix of cosine similarities between the TF-IDF representations of sentences in a document. Segmentation points are identified wherever there are sudden drops in the average similarity of sentences. The C99 algorithm enables us to coalesce individual sentences into larger, more meaningful segments that align more closely with coherent thoughts or ideas in the text. Consequently, we use these segmented units, rather than individual sentences, as the unit of analysis for our study.

3. Removing Boilerplate Segments

To identify and remove boilerplate segments in analyst reports, we employ a machine learning model specifically designed and trained for this purpose.

To construct the training data set for our model, we first identify the top 20 brokers producing the highest volume of reports each year. From each of those brokers in each sample year, we sample 1,000 of their reports. We then identify the top 10% most frequently repeated segments within those reports. For a segment to be classified as a positive example, it must satisfy two criteria: it is among the top 10% most frequently repeated segments, and it is repeated at least five times by the same broker within the same year.

Negative examples, or segments least likely to be boilerplates, are identified by randomly selecting 10 segments with no repetition in each broker-year sample. To ensure balance within our data set, we randomly sample from the remaining non-boilerplate segments to achieve a one-to-ten ratio of positive to negative examples. This results in a training data set of 547,790 examples, comprising 54,779 positive examples of boilerplate segments and 493,011 negative examples of non-boilerplate segments. The data set is split into training, validation, and testing sets, using an 80/10/10 ratio.

Our approach to identifying boilerplate segments in reports makes use of the SentenceTransformer model, specifically the all-mpnet-base-v2 variant. This model builds on an architecture similar to BERT, but focuses on creating high-quality sentence-level embeddings instead of token-level embeddings. This is particularly beneficial for our task as it views sentences and segments as distinct units of meaning, and thereby generates more effective and contextually relevant embeddings. To generate these embeddings, the SentenceTransformer model employs a mean pooling operation on the output of the transformer network, i.e., it creates a fixed-length sentence embedding by averaging all token embeddings. This operation gives us a representation of each sentence in a 768-dimensional vector space. For a segment containing multiple sentences, we compute the mean of all the sentence embeddings within that segment to yield a representative vector. This is an essential step because it allows us to convert segments of varying lengths into fixed-length representations, which can then be directly fed into our classification model. We find that this simple aggregation method performs well in capturing the overall semantic context of a segment.

BERT, with its bi-directional context understanding, is known to be effective for a broad range of NLP tasks (Devlin et al. 2018). A standard practice is to fine-tune the pre-trained BERT model on a specific task, which adjusts all the model parameters. This approach typically achieves high performance as it enables the pre-trained BERT model to learn from the specifics of the task, capitalizing on its general language understanding capabilities while adapting to task-specific nuances. However, given the context of our research, we adopt a different strategy that leverages both the representation power of the pre-trained BERT model and the efficiency of a classification head. Rather than fine-tuning and adjusting all parameters of the model, we freeze the parameters of the BERT model (i.e., the embeddings are fixed) and add a classification head to the model. The classification head takes embeddings from the BERT model as inputs and processes them with two hidden layers: the first layer contains 16 neurons, and the second layer contains 8 neurons. Each layer applies a Rectified Linear Unit (ReLU) activation to introduce non-linearity. Following the processing of the embeddings by the two layers, the resulting output vector is directed towards a softmax layer, which computes the probability of each segment as boilerplate or not.

The choice of the above strategy (architecture) is motivated by a number of considerations. First, our main generative AI model leverages retrieval augmented generation (RAG) that uses the all-mpnet-base-v2 embedding to help retrieve more context for culture-related information extraction. The chosen architecture allows us to maintain the consistency of embeddings across models. Second, we find that identifying boilerplate text is a relatively straightforward task that does not require the full-scale fine-tuning of the model, which would be computationally expensive and time-consuming. By freezing the parameters of the segment representation model and deploying only the classification head, we optimize computational efficiency and streamline the training process.

The trained classification model achieves good performance, with an Area Under the Curve (AUC) of 0.966 on the test set. The false positive rate is 0.093 and the false negative rate is 0.073.

Table IA1 in the Internet Appendix lists predicted boilerplate probabilities and boilerplate examples, sorted by decile. We retain segments with a boilerplate probability of 0.22 (the sample median) or lower.

4. Identifying Culture-related Segments

We identify segments related to corporate culture through a two-step procedure.

In step 1, we start with an exhaustive text search using two sets of keywords. The initial set of keywords is based on the word set explicitly about corporate culture, identifying a total of 5,541 relevant segments.[1] We also employ a second, more flexible set of keywords, which match all segments containing the word "culture(s)" or "cultural," excluding those already identified to avoid duplication. This second search results in a larger set of 46,795 segments. These segments contain potentially relevant mentions of corporate culture, although their meaning could be ambiguous. The word "culture(s)," and to a lesser extent the word "cultural," may refer to biological or social context.

To address these ambiguities, we employ generative AI for word sense disambiguation (WSD). Table IA2 Panel A shows the prompt. Our method matches the word "culture(s)" or "cultural" with one of the three definitions from dictionary.com.[2] The following examples illustrate our application of WSD:

> 1. "*Organization structure, talent model and deep bench of UNH make for a strong competitive advantage. The passion for excellence, humility, restlessness and desire to win is a culture that keeps UNH at the top of its game, and we expect will make it hard for fast-followers in the Large Cap MCO space and new Big Tech entrants such as AMZN and AAPL to catch up.*" -> Organizational
>
> 2. "*Demand for specialty proteins, probiotics, and cultures supported pricing gains. Continued strength in demand should contribute to a 1% and 4% YoY increase in sales for 1Q14 and 2014, respectively.*" -> Biological
>
> 3. "*Cultural hurdles more relative than price. There's no question that web conferencing is significantly less expensive than face-to-face meetings that require corporate travel (although it doesn't always replace a face-to-face meeting). To our knowledge, no one is questioning the value proposition of web conferencing.*" -> Societal

We exclude segments in which the discussion of culture is classified as in a biological or societal context, resulting in a final set of 41,038 segments.

It is possible that there are segments about corporate culture without mentioning explicit words or phrases. Consider the following segment "*One word we have heard from BBY's management team, a word that has led to the highest service levels at retailers and often the most successful ones is empowerment. From Wal-Mart in its heyday to Home Depot to Costco to Bed Bath and Beyond, empowering employees has been a critical element to success among retailers*." This segment, while not mentioning 'culture,' a human reader will conclude that it discusses a key aspect of corporate culture: employee empowerment.

In step 2, we fine-tune a BERT model to identify culture-related segments that lack specific keywords. The construction of our training set involves using segments, identified in step 1 as containing relevant keywords, as positive examples (culture = 1). Conversely, we include randomly selected segments without those keywords as negative examples (culture = 0). This training set is used to fine-tune the model, which is then deployed across all segments (excluding those identified in step 1). Based on the model's predictions, we sort these segments by percentile rankings of predicted probabilities. We focus on the top 5% of segments with the highest predicted probabilities of relating to culture.

Although the trained model achieves a high AUC (at 0.981), we observe a significant number of false positives. The segments predicted with high probabilities often pertain to other intangible aspects such as leadership or strategy, rather than corporate culture. To address this issue, we integrate the

[1] We use the following phrases for this exact matching process: "corporate culture," "company culture," "company's culture," "firm culture," "firm's culture," "organizational culture," "workplace culture," "business culture," and "culture in the company."

[2] For the segments containing the word "cultural," the definition is simply "Cultural: of or relating to culture, defined as …", followed by the corresponding definition for "culture."

capabilities of ChatGPT for an additional layer of filtering. Table IA2 Panel B shows the prompt. The prompt instructs ChatGPT to assess whether each input segment is relevant to corporate culture, based on any of the following four definitions of corporate or organizational culture: principles and values guiding employees (Guiso, Sapienza, and Zingales 2015), shared beliefs, assumptions, values, or preferences driving group behaviors (Li and Van den Steen 2021), norms and values widely shared and strongly held in the organization (O'Reilly and Chatman 1996), or an informal institution characterized by behavioral patterns reinforced by events, people, and systems (Grennan and Li 2023).

For each input segment, ChatGPT is asked to provide a brief explanation (not exceeding 50 words) justifying whether the segment discusses corporate or organizational culture topics as per the provided definitions, and a classification of the segment as either "Culture" or "No Culture." Only those segments classified as "Culture" are retained. This step adds 51,434 segments. Our final data set comprises 92,472 culture-related segments (41,038 segments from step 1 + 51,434 segments from step 2).

Figures IA 1-4 provide some overview of analysts' culture-related discussions in reports.

B. Matching Analyst Name in Reports to Analyst ID (AMASKCD) in I/B/E/S

We match lead analyst (i.e., the first author of a report) name to analyst ID (AMASKCD) in the I/B/E/S database as follows.

First, to unmask abbreviated broker names and analyst names from I/B/E/S, we manually search each broker's full name and its analysts from Capital IQ. Our matching process involves three steps: 1) we match abbreviated broker names in I/B/E/S (ESTIMID) to full broker names in Capital IQ by resemblance; 2) we ascertain the match in Step 1 by matching analyst names (ANALYST) in I/B/E/S with those in Capital IQ using the last name and first name initial; and 3) we supplement the above two steps by checking whether Capital IQ analysts' stock coverage is the same as that by matched I/B/E/S analysts. Of the 1,075 broker names in I/B/E/S, we are able to unmask full names for 928 brokers (an 86.3% matching rate).

We then obtain analyst information, including biography and prefix (Mr. versus Ms.), from their employment history in Capital IQ. In the end, we are able to unmask 13,164 out of the 14,909 analysts in the I/B/E/S Detail Recommendations file (an 88.3% matching rate).

Second, to match each analyst in the report sample to analyst ID (AMASKCD) in the I/B/E/S data set, we match each analyst's name in Investext to our unmasked broker names and analyst names in the I/B/E/S-Capital IQ merged sample as described above. Our matching proceeds as follows: 1) we match each broker in Investext to broker name and ID (EMASKCD) in the I/B/E/S-Capital IQ merged file; of the 1,006 unique brokers in Investext, we can link 443 brokers with EMASKCD – analysts affiliated with these 443 brokers produce 91% of the reports in our report sample; and 2) for cases in which Investext has lead analyst's full first name and full last name, we match analyst name in Investext to analyst name and ID (AMASKCD) in the I/B/E/S-Capital IQ merged file; we further verify this match if there is also a match between broker name and EMASKCD established above. In the end, we are able to uncover AMASKCD for 7,921 analysts, representing 78% of the analysts affiliated with the 443 brokers in our analyst report sample.

Our final sample comprises 1,744,540 reports covering 38,530 firm-year observations for 2,988 unique firms over the period 2000-2020.

C. Matching Analyst Reports to I/B/E/S Forecast Data

When examining whether and how analysts' views of culture impact price formation at the report level, we need to control for each report's quantitative and qualitative output (i.e., earnings forecast, stock recommendation, and target price). As a result, we need to match analyst reports from Investext in our sample with I/B/E/S forecast data following prior work (e.g., Huang, Zang, and Zheng 2014).

We employ a similar approach to link each report with its earnings forecast, stock recommendation, and target price in I/B/E/S. Here is an illustration of the process using earnings forecasts as an example.

Each report in our sample (from Section B) has a report date (DATE), a firm ID from I/B/E/S (CUSIP), an analyst ID from I/B/E/S (AMASKCD), and a broker ID from I/B/E/S (EMASKCD). Each earnings forecast in the I/B/E/S Detail file has an announcement date (ANNDATS), a review date (REVDATS), a firm ID (CUSIP), an analyst ID (AMASKCD), and a broker ID (EMASKCD). The announcement date is the day when an analyst revises her estimate and provides a forecast in a report. The review date is the day when an analyst confirms to I/B/E/S that her outstanding forecast is current. In I/B/E/S, a forecast is considered valid during the period from its announcement date until its review date. Within the period, analysts may issue multiple reports reiterating a forecast, but these reiterations have no separate entries in I/B/E/S.

We use a matching window, which is two days before a report's announcement date to two days after its review date, to match a report from Investext in our sample to an earnings forecast from I/B/E/S if the report date (from Investext) is within the "matching window," and there is a match of CUSIP-AMASKCD-EMASKCD between a report in our sample and an I/B/E/S earnings forecast.

Of the 1,744,540 reports in our sample (from Section B), we are able to match stock recommendations from I/B/E/S for 1,413,260 reports (an 81.0% matching rate), representing 36,100 firm-year observations associated with 2,898 unique firms; we are able to match target prices for 1,402,233 reports (an 80.4% matching rate), representing 35,673 firm-year observations and 2,887 unique firms. The samples used in Table 8 are smaller due to data availability for control variables.

Finally, we are able to match earnings estimates, stock recommendations, and target prices from I/B/E/S for 1,089,760 reports (a 62.5% matching rate), representing 34,314 firm-year observations associated with 2,858 firms. The sample used in Table 9 is smaller due to data availability for control variables.

Table IA1
Examples of the boilerplate segments

The table provides examples of the boilerplate segments in analyst reports, sorted by the predicted probability of a segment being boilerplate using a fine-tuned BERT model (in descending order). We retain segments with the predicted probability at 0.22 (the sample median) or lower.

Decile	Probability	Example
10	0.998	The investments or services contained or referred to in this report may not be suitable for you and it is recommended that you consult an independent investment advisor if you are in doubt about such investments or investment services. Nothing in this report constitutes investment, legal, accounting or tax advice or a representation that any investment or strategy is suitable or appropriate to your individual circumstances or otherwise constitutes a personal recommendation to you. CS does not offer advice on the tax consequences of investment and you are advised to contact an independent tax adviser.
9	0.979	EEA -The securities and related financial instruments described herein may not be eligible for sale in all jurisdictions or to certain categories of investors.
8	0.925	Distribution of ratings: See the distribution of ratings disclosure above. Price Chart: See the price chart, with changes of ratings and price targets in prior periods, above, or, if electronic format or if with respect to multiple companies which are the subject of this report, on the DBSI website at http://gm.db.com.
7	0.871	Suspended -the company rating, target price and earnings estimates have been temporarily suspended. For disclosure purposes, Evercore Group's prior “Overweight,” “Equal-Weight” and “Underweight” ratings were viewed as “Buy,” “Hold” and “Sell,” respectively. Evercore ISI utilizes an alternate rating system for companies covered by analysts who use a model portfolio-based approach to determine a company's investment recommendation.
6	0.858	As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of the report and investors should consider this report as only a single factor in making their investment decision.
5	0.264	He also noted the Fed now has much more robust tools to guard against systemic risk vs. overseeing individual institutions. There were no questions about the history of the supervision process of large bank holding companies within the Atlanta Fed (district six) or more generally the troubled Georgia and Florida real estate markets.
4	0.054	These contracts usually have 1-5 year terms. Since residential customers generally fall under local government jurisdictions, the contracts are negotiated with the municipality -which in turn, will bill the residential customer through taxes. For larger commercial and industrial customers, the company will negotiate directly with the end-user.
3	0.028	Integration of recent acquisitions. Our downside risk for the shares is $10, which represents just over 10x our 2012 EPS estimate of $0.95 and is typically the bottom end of the company's historical multiple range.
2	0.003	We include these gains in our calculation of ongoing earnings because they are not entirely one-time though they are infrequent items. We have raised our 2005 EPS estimate to $1.50 from $1.40 to include an expected $0.11 fourth quarter investment gain on the sale of an interest in a coal-fired power plant in Georgia.
1	0.000	“Harman is still in a transitional phase with its highly anticipated scalable infotainment backlog not launching until fiscal 2013/2014. And while we acknowledge the company has made significant progress on the cost side, Harman will have to consistently execute on those cost cutting initiatives for the next several quarters to help prop-up its low-price and low-margin customized business.”

Table IA2
Prompts to filter culture-related segments in analyst reports

This table presents detailed instructions given to generative AI models to filter corporate culture-related segments in analyst reports. Panel A shows the prompt for word sense disambiguation relating to the words "culture(s)" and "cultural." Panel B shows the prompt to determine whether a segment with a high predicted probability of relating to culture is really about corporate culture.

Panel A: Prompt to perform word sense disambiguation

```
Perform word sense disambiguation on the word 'culture(s)' or 'cultural' in the input segments.
These segments are from sell-side equity analyst research reports on companies. Classify each input
segment into one of the three categories: 'organizational,' 'societal,' or 'biological.' Definitions
for these categories are as follows:

* organizational culture: 'The values, typical practices, and goals of a business or other
organization, especially a large corporation. e.g., Their corporate "culture" frowns on avoiding
risk. We recognize a cost-cutting "culture". '
* societal culture: 'The behaviors and beliefs characteristic of a particular group of people, as a
social, ethnic, professional, or age group (usually used in combination). e.g., the youth "culture";
the drug "culture".
* biological culture: 'Biology. the cultivation of microorganisms, as bacteria, or of tissues, for
scientific study, medicinal use, etc. or the product or growth resulting from such cultivation.
e.g., cell "culture".'

For each input segment, provide:
    1. The input ID (input_id).
    2. The classification of the word 'culture(s)' or 'cultural' in input text as 'organizational,'
'societal,' or 'biological.'

Format your response in JSON. Make sure you process all of the inputs. Response format:

{"all results":
    [
        {
            "input_id": "XXXX",
            "classification": "organizational/societal/biological",
        },
        {
            "input_id": "YYYY",
            "classification": "organizational/societal/biological",
        },
        ... (other inputs)
    ]
}
```

Panel B: Prompt to determine if a segment is really about corporate culture

```
Assess whether each of the following input segments is relevant to corporate or organizational
culture. The segments are from sell-side equity analyst research reports on companies. An input
segment is relevant to corporate or organizational culture if it discusses topics consistent with
any of the following definitions for corporate or organizational culture:

* "principles and values that should inform the behavior of all the firms’ employees."
* "a group’s shared beliefs, assumptions, values, or preferences that then drive that group’s
behaviors."
* "a set of norms and values that are widely shared and strongly held throughout the organization."
* "an informal institution typified by patterns of behavior and reinforced by events, people, and
systems."

For each input segment, provide:
    1. The input ID (input_id).
    2. Brief reasons (in 50 words or less) that explain if the segment contains discussions about
corporate or organizational culture topics using ANY of the definitions provided above.
    3. Classification of the segment as 'Culture' if it is relevant to corporate or organizational
culture, or as 'No Culture' if it does not.

Format your response in JSON. Make sure you process all of the inputs. Response format:

{"all results":
    [
        {
            "input_id": "XXXX",
```

```
            "explanation": "This segment is relevant to corporate/organizational culture because...",
            "classification": "Culture",
        },
        {
            "input_id": "YYYY",
            "explanation": "This segment is not relevant to corporate/organizational culture because it mainly discusses...",
            "classification": "No Culture",
        },
        ... (other segments)
    ]
}
```

Table IA3
A list of cultural values in the literature

This table lists cultural values examined by prior work in the literature. The last column lists the nine cultural values that analysts refer to when discussing culture in their reports extracted by generative AI in our paper.

Guiso, Sapienza, and Zingales (2015)	Grennan (2019)	Li et al. (2021)	Graham et al. (2022a, 2022b)	Our paper
(1)	(2)	(3)	(4)	(5)
Communication	Adaptability	Innovation	Adaptability	Adaptability
Community	Collaboration	Integrity	Collaboration	Customer-oriented
Hard work	Customer-orientation	Quality	Community	Innovation
Innovation	Detail-orientation	Respect	Customer-orientation	Integrity
Integrity	Integrity	Teamwork	Detail-orientation	Operations-oriented
Quality	Results-orientation		Integrity	People-oriented
Respect	Transparency		Results-orientation	Results-oriented
Safety				Risk control
Teamwork				Teamwork

Table IA4
Representative examples of the extracted cultural values, their causes, and their effects

The table provides some representative examples of the extracted cultural values, and their causes and effects. The causes are grouped into three categories: events, people, and systems suggested by Guiso, Sapienza, and Zingales (2015), Graham et al. (2022a, 2022b), and Grennan and Li (2023).

Panel A: Different cultural values

Value	Example
Adaptability	adaptive culture, adaptive corporate culture, change-oriented culture, resilient culture, proactive culture, continuous improvement culture, evolving culture, transformative culture, transformational culture, agile culture
Customer-oriented	customer-centric culture, sales-driven culture, sales-oriented culture, service-oriented culture, customer-focused culture, customer-centric, client-focused culture, client-centric culture, consumer-centric culture, customer-obsessed culture
Innovation	innovative culture, entrepreneurial culture, growth-oriented culture, innovative corporate culture, data-driven culture, technology-driven culture, innovation-driven culture, knowledge-driven culture, creative and innovative culture, entrepreneurial and decentralized culture
Integrity	accountable culture, community-oriented culture, ethical corporate culture, socially responsible culture, accountability culture, accountability-driven culture, values-driven culture, integrity-based culture, transparent culture, integrity-driven culture
Miscellaneous	challenging corporate culture, acquisitive culture, ambitious culture, stable corporate culture, dedicated corporate culture, experienced corporate culture, focused culture, unique culture, traditional corporate culture, long-term focused culture
Operations-oriented	decentralized culture, cost-conscious culture, efficiency-driven culture, efficiency-oriented culture, quality-focused culture, cost-cutting culture, detail-oriented culture, centralized culture, process-oriented culture, operational culture
People-oriented	employee-centric culture, inclusive culture, diverse corporate culture, family-oriented culture, people-centric culture, talent-focused culture, people-focused culture, empowering culture, internal-promotion culture, supportive culture
Results-oriented	performance-driven culture, results-oriented culture, competitive culture, aggressive culture, profit-driven culture, goal-oriented culture, high-performance culture, shareholder-focused culture, winning culture, success-driven culture
Risk control	disciplined culture, conservative culture, risk-averse culture, cautious culture, risk-aware culture, safety-oriented culture, compliance-oriented culture, financially disciplined culture, prudent culture, risk management culture
Teamwork	collaborative culture, integration-oriented culture, team-oriented culture, cohesive corporate culture, partnership-oriented culture, cooperative culture, team-based culture, silos culture, alignment-oriented culture, collegial culture

Panel B: Different causes of culture

Category	Cause	Example
Event	COVID-19	covid-19 pandemic, disruption from covid-19, response to covid-19 pandemic, uncertainty related to the covid-19 pandemic, covid disruptions
Event	Disruptive technology	disruptive products, next-generation technology, focus on disruptive innovation, phase of disruptive technology, breakthrough innovations through AI
Event	Economic downturn	financial crisis, economic cycle, economic downturn, economic environment deterioration, weak macro environment and jobs market
Event	Internal conflicts	Bureaucracy, internal politics, disconnect between it department and senior management, history of frequent strike activity, internal power struggles
Event	Management change	leadership change, new management team, management turnover, new CEO, separation of chairman and CEO roles
Event	Market expansion	international expansion, expansion into new markets, global presence, rapid expansion, seeking new areas of growth

Event	Mergers and acquisitions	acquisitions, strategic acquisitions, mergers and acquisitions, M&A activities, integration of acquired businesses
Event	Regulatory issues	regulatory actions, challenging regulatory environments, unknowns within consumer regulatory agency, regulatory changes, SEC investigation
Event	Shareholder activism	shareholder pressure, significant insider ownership, interaction with activist investors, agreement with activist investor for board changes, proxy fight with activist investor
Event	Strategic transformation	organizational restructuring, strategic initiatives, reorganization, transition to solutions-based focus, strategic changes, business transformation
People	Customer relations	focus on customer service, best-of-breed customer service, deep client relationships, decades of high-quality service, direct relationships with end-users
People	Management team	experienced management team, strong management team, CEO's leadership, visionary leadership, long-tenured management team
System	Business relationship	local management with deep roots in each community, valuable commercial client relationships, establishment of unique relationship with independent agents, ability to attract and foster close and long-lasting business relationships, strategic partnerships
System	Business strategy	training salesforce in value-over-volume strategy, focus on cost management, differentiated merchandising strategy, management's aggressive expansion initiative
System	Compensation structure	compensation structure, competitive compensation programs, compensation structure emphasizing incentive pay, incentive compensation structure, employee stock ownership
System	Employee hiring and retention	promotion from within, extensive training programs, long tenure of employees, workforce reduction, resisting layoffs during recession
System	Workplace safety	desire to minimize personal risk for employees, industry's effort to improve safety practices, independent safety oversight committee, efforts to improve safety practices, fair hearing and remedy process for workers' grievances
	Miscellaneous	structural attributes, resources, competitive environment, deep and broad industry expertise, accounting practices

Panel C: Different effects of culture

Consequence	**Example**
Business relationship	strong customer relationships, cross-selling opportunities, critical industry relationships, retaining valuable commercial client relationships, stronger franchisee alignment
Customer satisfaction	improved customer service, improved customer experience, customer satisfaction, customer loyalty, improved customer satisfaction
Diversity, equity, and inclusion	improved diversity of leadership team, development of a diverse talent base, lack of diversity in board composition, promotion of more women, toxic culture of sexual harassment
Employee satisfaction	employee turnover, employee retention, low employee turnover, employee satisfaction, employee ownership
ESG practices	corporate governance weaknesses, environmental sustainability efforts, esg practices, enhanced governance practices, development of environmentally and ethically responsible products
Innovation	accelerated development of desirable new products, product innovation, new product development, technological leadership, focus on innovation
Internal conflicts	potential for business conflicts, wrestling with production planning, resistance to change, potential muddled strategy and infighting, management distraction
Investor relations	attractiveness to investors, rebuilding investor confidence, alignment of management and shareholder interests, shareholder friendliness, improved communication with investors
Management change and retention	loss of key personnel, management resource strain, smooth leadership transition, strong management team, management turnover

Market share and growth	revenue growth, market share gains, increased market share, expansion into new markets, establishing a strong presence in key markets around the world
Mergers and acquisitions	successful integration of acquisitions, successful acquisitions, challenges in integration, lower-than-expected synergies, M&A strategy
Misconduct	management protecting their own interests over investors, unusual accounting moves, massive legal liabilities, multiple scandals, legal troubles
Miscellaneous	Competitive advantage, long-term success, improved operations, continued success, positive geographic mix
Profitability	margin expansion, improved profitability, cost savings, increased profitability, more stable levels of profitability
Resilience	business resilience, resilience in the next downturn, resilience during recession, successful weathering of recent market volatility and macroeconomic uncertainty, persistent corporate momentum
Risk management	focus on risk management, focus and importance placed on risk management, handling credit risk well, improved credit quality, minimized franchise risk
Shareholder value	consistently above-average returns, strong balance sheet, returns exceeding cost of capital for longer periods, enhanced shareholder value, alignment of interests with shareholders

Table IA5
Examples of the extracted cause or effect relations

This table provides examples of the extracted cause or effect relations by generative AI. In each example, a snippet of the culture-related segment is provided, and the extracted terms with corresponding canonicalized terms in parentheses and the cause or effect relations are highlighted in boldface.

Example 1: Coca-Cola's CEO Neville Isdell presented for the company at the CAGNY conference in Arizona this morning. While there was not much new news in the presentation, the company's tone has changed meaningfully from Isdell's presentation at CAGNY two years ago (when KO was promising the market very little) to today, when the company has greater confidence that its long-term algorithm is both working and sustainable. We highlight what we think were a few of the key takeaways below:
• The company highlighted the improving performance in 2006 as KO has moved through its Manifesto For Growth strategy and now enters a phase of likely sustainable growth with a focus on growing the core brands, capturing emerging platforms by establishing a culture of innovation, and providing franchise leadership to the bottlers (which was consistent with what we heard from CCE yesterday, with a greater focus on increasing collaboration with the bottling system).
• KO is entering 2007 with some of the strongest growth momentum in the last several years, and the company has demonstrated that they can continue to post solid growth despite underperformance in key markets.NA should remain a challenge in 2007, but KO has clearly shown an ability to turn around problem markets and seems very comfortable with its long-term growth model given strength in the balance of the business. While it's not yet clear what stage in the turnaround we have entered for key markets such as Japan, Germany, and India, KO expressed confidence in the recent improvements in these areas and we are encouraged by the results…

Extracted relation(s):

- **Manifesto for growth strategy (Business strategy) → Innovative culture (Innovation).**
- **Innovative culture (Innovation) → Strongest growth momentum in the last several years (Market share and growth).**
- **Innovative culture (Innovation) → Improving performance in 2006 (Profitability).**

Ref: This report was written by John A. Faucher from JP Morgan for Coca-Cola Co. released on 2/22/2007.

Example 2: …WGL's strategic objectives to improve its Gas Distribution business so that it provides stable and growing revenue, foster a high-performance culture throughout the organization, and develop profitable retail energy-related businesses should become achievable as the impact of its outsourcing agreement becomes more visible.…

Extracted relation(s):

- **Strategic objectives (Business strategy) → High-performance culture (Results-oriented culture).**

Ref: This report was written by Joanne M. Fairechio from Janney Montgomery Scott LLC for WGL Holdings Inc. released on 11/27/2007.

Example 3: Our $46 Dec-12 Price Target is based on 12.0x our 2013 EPS of $3.85 (above M's trailing 3-year 10.5x avg, but 230bps below its dept store peer average of 14.3x). Focused on company specific initiatives (My Macy's, Magic Selling, and Direct/Omni-channel) M has separated itself from moderate peers (JCP/KSS) executing on a three-tiered strategy (brands, fashion, price). With the turn in the selling culture taking place just last summer (according to CEO Lundgren) and with gross margin drivers on the horizon (Omnichannel and price optimization) we see double digit earnings growth through 2015.

Extracted relation(s):

- **Company specific initiatives (my Macy's, magic selling, direct/omni-channel) (Business strategy) → Selling culture with a focus on brands, fashion, and price (Customer-oriented).**

- **Selling culture with a focus on brands, fashion, and price (Customer-oriented) → Double digit earnings growth through 2015 (Market share and growth).**

Ref: This report was written by Matthew R. Boss from JP Morgan for Macy's Inc. released on 3/12/2012.

Example 4: It was championed by Starbucks' chairman Howard Schultz since early in Starbucks' existence. Health care costs have increased dramatically over the past several years, and continue to pressure the company's operating margins. However, we do not expect Starbucks to move away from the health benefit program, as it is an important part of its culture and its ability to attract good employees...

Extracted relation(s):

- **Howard Schultz (Management team) → Employee-centric culture (People-oriented).**
- **Employee-centric culture (People-oriented) → Ability to attract good employees (Employee satisfaction).**

Ref: This report was written by Ashley R. Woodruff from Bear, Stearns & Co., Inc. for Starbucks Corporation released on 5/18/2005.

Example 5: … ARG has a 20% compound annual return to shareowners in its 20 years as a public company, on strong growth, led by a scrappy entrepreneurial culture bred by its founder and CEO, who maintains a 10% stake in the company…

Extracted relation(s):

- **Founder's leadership (Management team) → Entrepreneurial culture (Innovation).**
- **Entrepreneurial culture (Innovation) → 20% compound annual return to shareowners (Shareholder value).**

Ref: This report was written by Edward Hoonshik Yang from CIBC Capital Markets Corp. for Airgas Inc released on 11/28/2007.

Example 6: Schein's U.S. Medical division appears poised for better-than-market revenue and earnings growth. Recent management changes have taken hold and have created a stronger, more team-oriented culture in the division. Management's improvements are bearing out in the favorable top-line performance of this business, which ranks among the strongest revenue growth rates in all of medical/surgical distribution. We believe that Schein's better-than-market revenue growth in U.S. Medical reflects two elements of this business: Schein's direct-mail orientation creates cost-advantages and its catalog-based revenue (which accounts for about one-half of U.S. Medical total revenue) is growing at 2 times the rate of salesperson-based revenue.

Extracted relation(s):

- **Team-oriented culture (Teamwork) → Better-than-market revenue growth (Market share and growth).**
- **Recent management changes (Management change) →Team-oriented culture (Teamwork).**

Ref: This report was written by David G. O'Neill from William Blair & Company for Henry Schein In. released on 6/20/2000.

Example 7: … XPO has also instilled a strong performance-based culture (salary plus incentive compensation tied to gross margin dollars and gross margin dollars per load), which we believe has contributed to the recent increase in productivity and should help drive revenue growth as the salesforce continues to mature (average tenure is a little over one year).We believe XPO is positioned well to gain market share as it has focused on high customer service while broadening its solution set-something few competitors have been able to achieve successfully. …

Extracted relation(s):

- **Performance-based culture (Results-oriented) → Potential revenue growth (Market share and growth).**

Ref: This report was written by Nathan Brochmann from William Blair & Company for XPO Inc. released on 6/19/2015.

Example 8: … Demand Remains Solid; Raising FVE to $88 17 Feb 2017. Arista reported strong results in its fourth quarter, with revenue increasing above our expectations. We are impressed by another year of stellar revenue growth, as the company's strategic focus on large customers' needs and its culture of product innovation are paying off. …

Extracted relation(s):

- **Innovative culture (Innovation) → Stellar revenue growth (Market share and growth).**
- **Strategic focus on large customers' needs (Customer relations) → Innovative culture (Innovation).**

Ref: This report was written by Ilya Kundozerov from Morningstar Inc. for Arista Networks Inc. released on 5/8/2017.

Example 9: Citigroup's management team has created a culture that is bottom line focused, where revenue and expenses are given equal weight in measuring success. Accordingly, Citigroup has a culture that is almost paranoid about keeping costs low and finding new methods to improve efficiency. In the multiple conversations we had with management, we found that there is the unspoken expectation that margins will improve from one year to the next. If per unit costs are not falling, management indicated that something is wrong. Scale also plays an important role in Citigroup's higher margins. Because Citigroup is larger than its peers, it has scale in every business in which it participates, so costs are expected to be low and margins high.

Extracted relation(s):

- **Efficiency-driven culture (Operations-oriented) → High margins (Profitability).**
- **Management's focus on revenue and expenses (Business strategy)→ Efficiency-driven culture (Operations-oriented).**

Ref: This report was written by Ken Worthington from CIBC Capital Market Corp. for Citigroup Inc. released on 10/23/2002.

Example 10: … BBY has had 32 customer centricity stores in place for the past 10 months, which strive to drive incremental sales growth by more closely matching the stores to the needs of the customers in the relevant communities. The 32 test stores focus on profitable customer segments, tailor the product offering to the target market, and build an operating model to meet the needs of the customer base. In addition, these stores create an owner-operator culture at the store level as decisions are made closer to the customer. The 32 test stores led BBY's comparable store sales in 4Q04 and also earned higher gross margins…

Extracted relation(s):

- **Customer-centric culture (Customer-oriented) → Earning higher gross margins (Profitability).**
- **Customer-centric culture (Customer-oriented) → Leading comparable store sales (Market share and growth).**

Ref: This report was written by Dana L. Telsey from Bear, Stearns & Co., Inc. for Best Buy Co Inc. released on 4/1/2004.

Table IA6
Correlation matrices for the firm-year and firm-analyst-year samples

This table presents the correlation matrices for samples used in different regression analyses. Panel A presents the correlations for variables at the firm-year level. Panel B presents the correlations for variables at the firm-analyst-year level. Definitions of the variables are provided in the Appendix. Superscripts a, b, and c indicate significance at the 1, 5, and 10 percent levels, respectively.

Panel A: The correlation matrix for the firm-year sample

Variable	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1. Culture discussion	1.000														
2. Firm size	0.280^{a}	1.000													
3. Ln(Firm age + 1)	0.045^{a}	0.323^{a}	1.000												
4. Sales growth	0.020	-0.068^{a}	-0.188^{a}	1.000											
5. ROA	0.111^{a}	0.070^{a}	0.052^{a}	0.154^{a}	1.000										
6. Leverage	0.011^{b}	0.248^{a}	0.101^{a}	-0.061^{a}	-0.130^{a}	1.000									
7. Tangibility	-0.046^{a}	0.156^{a}	0.193^{a}	-0.100^{a}	-0.048	0.256^{a}	1.000								
8. ROA volatility	-0.126^{a}	-0.263^{a}	-0.129^{a}	0.031^{c}	-0.391^{a}	0.003^{c}	0.008^{a}	1.000							
9. Large institutional ownership	-0.125^{a}	-0.242^{a}	-0.059^{a}	-0.009^{a}	-0.044	0.029^{b}	-0.102^{a}	-0.013^{b}	1.000						
10. Board independence	-0.217^{a}	-0.389^{a}	-0.055^{a}	0.022^{a}	-0.139^{a}	-0.013^{a}	-0.030^{a}	0.041	0.238^{a}	1.000					
11. CEO duality	0.046^{a}	0.116^{a}	0.079^{a}	-0.013	0.045^{a}	-0.019^{c}	0.041^{a}	-0.106^{a}	-0.068^{a}	-0.078^{a}	1.000				
12. Ln(Number of key people changes + 1)	0.168^{a}	0.318^{a}	0.126^{a}	-0.070^{a}	-0.081^{a}	0.033^{a}	-0.003^{a}	-0.012	-0.100^{b}	-0.070^{a}	-0.052^{a}	1.000			
13. Ln(Number of M&As + 1)	0.136^{a}	0.374^{a}	0.075^{a}	0.023^{a}	0.002	0.106^{a}	-0.035^{a}	-0.053^{a}	-0.149^{a}	-0.185^{a}	0.031^{a}	0.187^{a}	1.000		
14. Corporate culture	0.119^{a}	-0.145^{a}	-0.211^{a}	0.087^{a}	-0.018^{a}	-0.129^{a}	-0.227^{a}	0.085^{a}	0.042^{a}	-0.033^{a}	-0.075^{a}	0.103^{a}	0.014	1.000	
15. Employee culture rating	0.050^{a}	0.054^{a}	0.064^{a}	-0.011	0.010	0.113^{a}	0.003	-0.058^{a}	0.201^{a}	0.132^{a}	-0.088^{a}	-0.008	-0.022^{b}	0.157^{a}	1.000
16. Ln(Number of employee reviews + 1)	0.298^{a}	0.467^{a}	0.159^{a}	-0.072^{a}	0.107^{a}	0.121^{a}	-0.054^{a}	-0.153^{a}	0.015^{c}	-0.200^{a}	0.000	0.256^{a}	0.196^{a}	0.272^{a}	0.458^{a}

Panel B: The correlation matrix for the firm-analyst-year sample

Variable	1	2	3	4	5	6	7	8	9
1. Culture discussion	1.000								
2. Star analyst	0.016[a]	1.000							
3. Female	0.025[a]	-0.032[a]	1.000						
4. Forecast horizon	-0.023[a]	0.049[a]	-0.003	1.000					
5. General experience	0.035[a]	0.108[a]	-0.030[a]	-0.030[a]	1.000				
6. Firm experience	0.048[a]	0.113[a]	-0.032[a]	0.013[a]	0.590[a]	1.000			
7. Number of industries followed	0.021[a]	0.032[a]	-0.022[a]	-0.113[a]	0.154[a]	0.068[a]	1.000		
8. Number of firms followed	-0.003	0.116[a]	-0.060[a]	0.018[a]	0.307[a]	0.180[a]	0.371[a]	1.000	
9. Forecast frequency	-0.003	0.098[a]	0.011[a]	-0.024[a]	0.039[a]	0.141[a]	-0.055[a]	0.089[a]	1.000
10. Broker size	0.049[a]	0.226[a]	0.035[a]	0.063[a]	-0.033[a]	0.012[a]	-0.085[a]	0.083[a]	0.104[a]

Figure IA1
Temporal trends in culture-related segments in analyst reports

The sample consists of 92,472 culture-related segments in analyst reports over the period 2000–2020.

Panel A: Temporal trends in culture-related reports and segments

# of reports

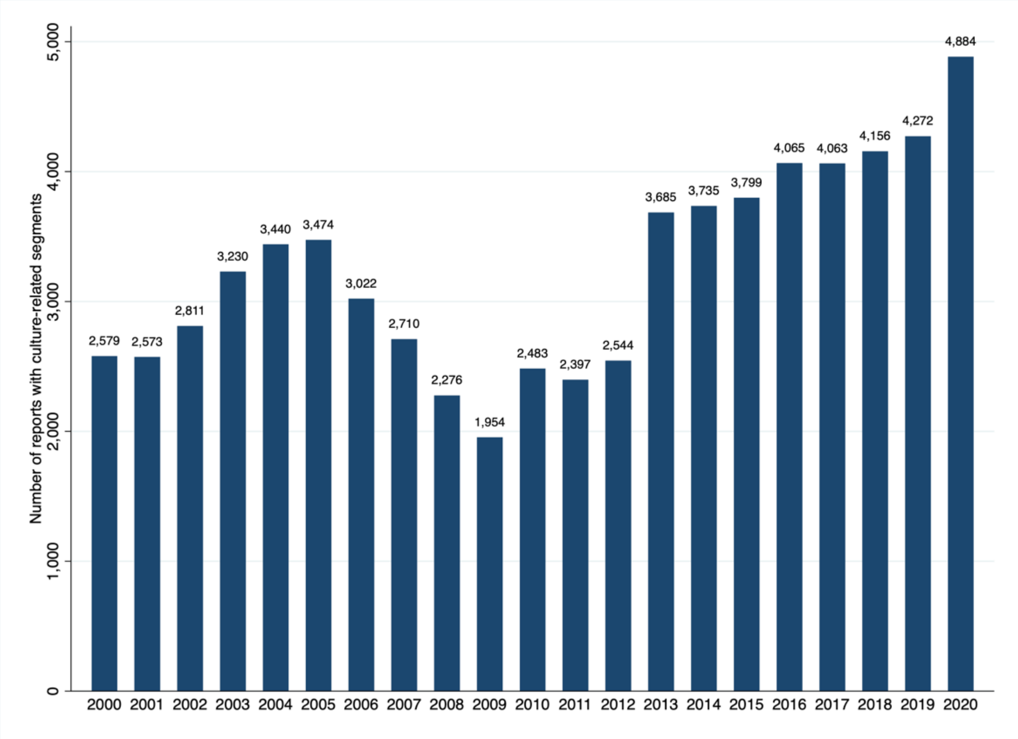

# of segments

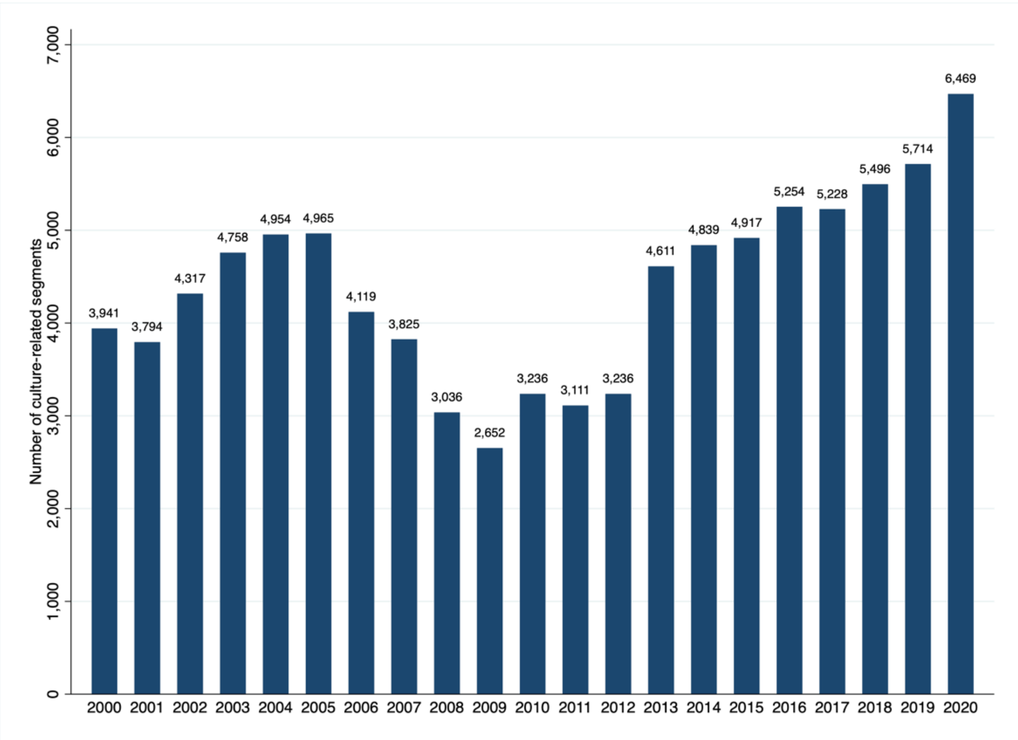

Panel B: Temporal trends in culture-related segments by cultural value

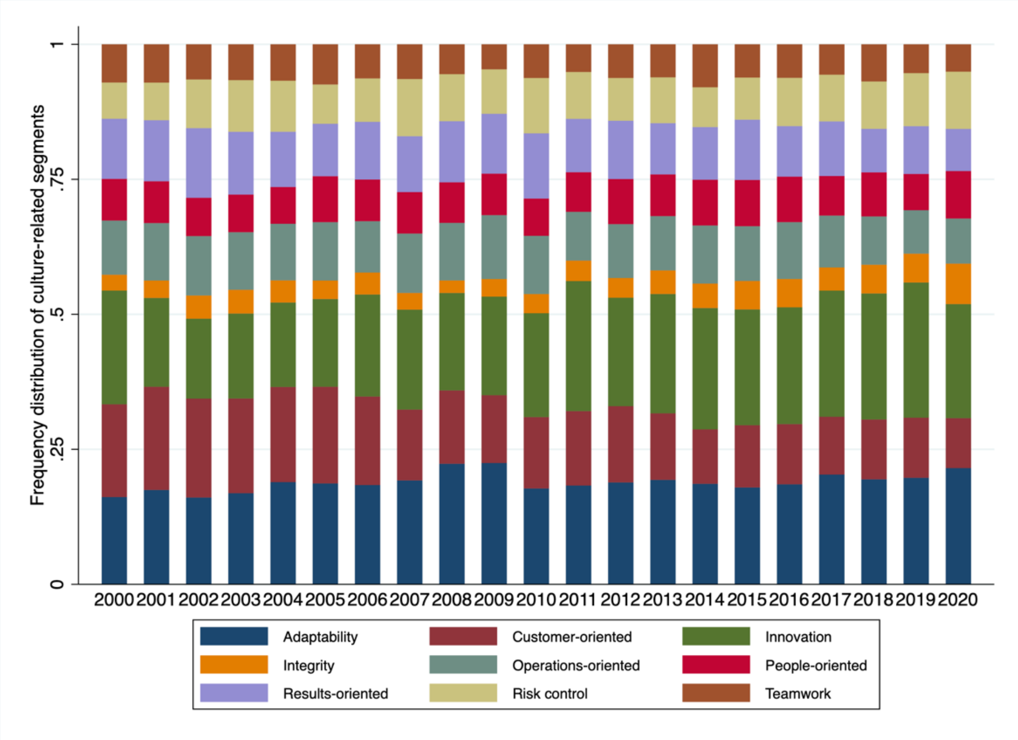

Figure IA2
Industry distributions of culture-related segments in analyst reports

The sample consists of 92,472 culture-related segments in analyst reports over the period 2000–2020. Our industry classification is based on Fama-French 12 industries.

Panel A: Industry distribution of culture-related segments in analyst reports

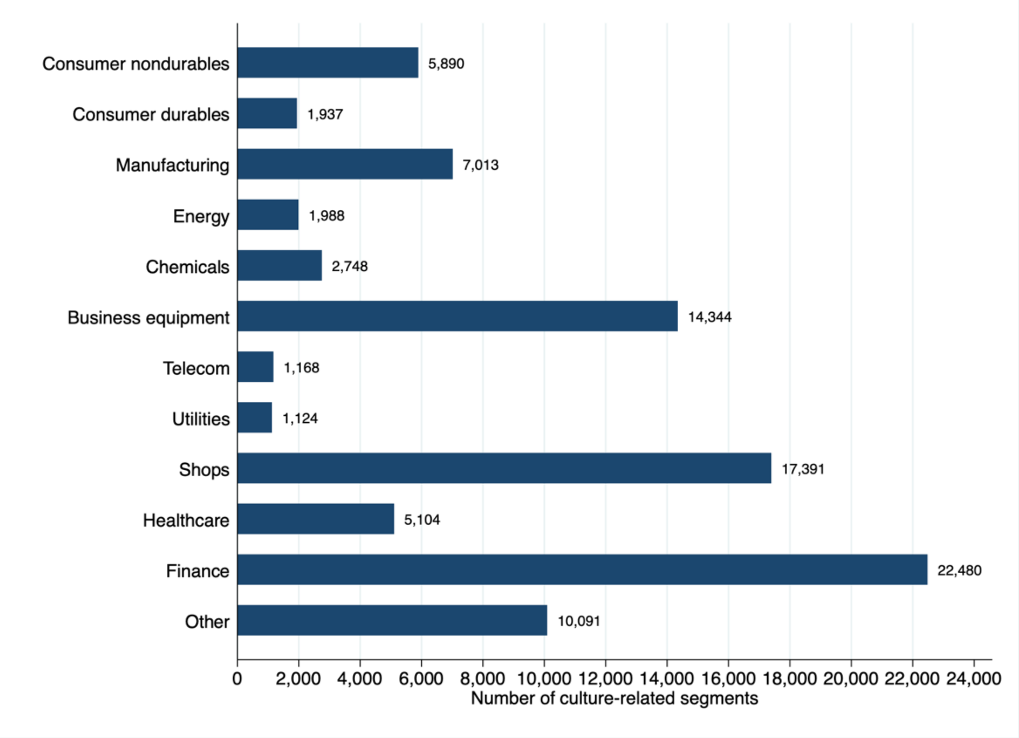

Panel B: Industry distribution of culture-related segments in analyst reports by cultural value

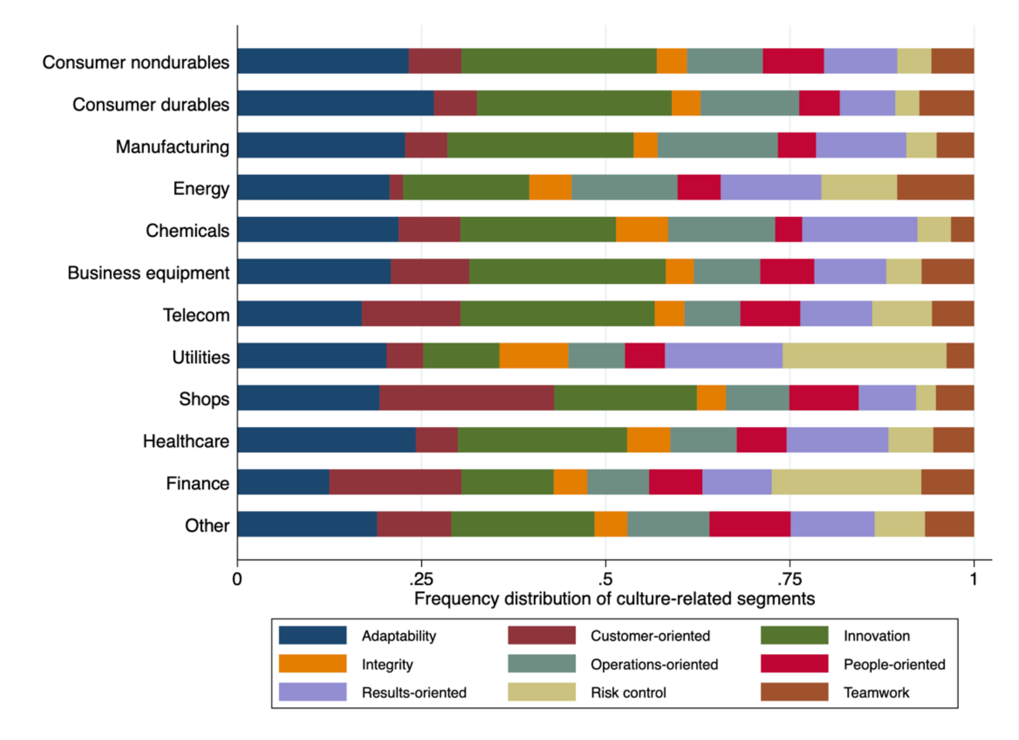

Figure IA3
Tone distribution of culture-related segments in analyst reports by cultural value

The sample consists of 92,472 culture-related segments in analyst reports over the period 2000–2020.