

Hidden in plain sight: Unraveling compensation disclosure bloat with generative AI and its impact on executive compensation

Lizi Burduli, Stephan Kramer

Received 27 August 2025 | Accepted 28 January 2026 | Published 21 April 2026

Abstract

Whether compensation contract design reflects efficient contracting or rent extraction is an ongoing debate in academic research and public discourse. We contribute to this debate by examining whether textual bloat in compensation contract disclosures is associated with excess CEO compensation. We construct a measure of bloat, defined as irrelevant, boilerplate, and redundant content, by summarizing firms' Compensation Discussion and Analysis sections with a large language model for a sample of S&P 1500 firms during 2011–2018. In line with our hypotheses, we find a positive association between bloat and excess CEO compensation. We find no empirical evidence that governance characteristics explain the magnitude of bloat in firms' compensation disclosures. Our findings suggest that bloated disclosures can be used as an instrument to obscure compensation levels that are unrelated to the economics of the firm.

Relevance to practice

As reflected in the EU Shareholder Rights Directive (Directive 2007/36/EC), European regulators seek to increase transparency and shareholder involvement, improve the oversight of directors' remuneration, and facilitate the flow of information. The findings of this study support these objectives by demonstrating that the substance and understandability of compensation disclosures, rather than their length, matter more for effective monitoring. In the Netherlands' stakeholder-oriented governance model, generative AI tools can help boards, auditors, and investors in evaluating whether compensation disclosures genuinely communicate important information.

Keywords

Generative AI, LLM, executive compensation, bloat, corporate governance

1. Introduction

The widening pay gap between CEOs and workers in recent decades underscores a central debate in corporate governance research: whether executive compensation reflects efficient contracting or managerial rent extraction. The efficient contracting view is based on classical agency theory (Jensen and Meckling 1976) and posits that boards design pay packages to align executive and shareholder interests. Conversely, the managerial power or rent extraction view (Bebchuk and Fried 2003) argues that powerful executives can leverage their influ-

ence over the board to extract rents, i.e., receive compensation packages that exceed what is justified by their talent or the firm's economics.

One potential mechanism to achieve this is obfuscation through bloat, i.e., adding redundant, overly complex, and irrelevant information to disclosures about compensation. In the United States, the Compensation Discussion and Analysis (CD&A) section of the proxy statement is the primary source of investors' information about executive compensation packages.

Prior research has attempted to capture obfuscation in other disclosure contexts by using traditional textual analysis methods, such as readability or sentiment (e.g., Li (2008); Loughran and McDonald (2014)). While these studies provide valuable first insights about the textual accessibility of disclosures, they are not intended to measure contextual irrelevance.

Recent advancements in generative Artificial Intelligence (AI) and Large Language Models (LLMs) offer an innovative methodological solution to this problem. Unlike older models, LLMs are pre-trained on vast datasets, enabling them to grasp the contextual significance and filter out irrelevant content in a manner that approximates human-like judgment. This allows for a novel and more precise measure of obfuscation. By tasking an LLM with summarizing a document to its essential core, we can quantify bloat as the proportion of the original text that is discarded as irrelevant or redundant. To date, no study has applied an LLM-based bloat measure to compensation disclosures to examine its link with excess pay. We address this gap by answering the following research question:

How can a measure of compensation disclosure bloat be developed using generative AI, and what is its relationship with excess executive compensation?

Using 7,786 firm-year observations from S&P 1500 companies between 2011 and 2018, we use a Large Language Model to construct summaries of CD&A sections and measure bloat as the difference between the length of the original document and its summary, scaled over the length of the original document. We quantify excess compensation using the methodology of Core et al. (2008) and conduct robustness checks using the approach by Larcker et al. (2011).

Our results show a significant positive association between bloat and excess compensation. This association remains robust after controlling for traditional textual metrics, which suggests our measure provides an incrementally informative dimension to the study of compensation design. We find no empirical evidence that governance mechanisms, as proxied by board gender diversity percentage and the CEO serving as chair of the board, are significant determinants of bloat. Taken at face value, this opposes the idea that firms that award the CEO with abnormally high pay packages and are poorly governed should have more bloated compensation disclosures (Core et al. 1999; Basu et al. 2007). However, another potential explanation is that both proxies used in this study are relatively invariant over time and therefore do not possess much explanatory power in our empirical design with a tight fixed effects structure, and that more research is needed to shed light on this relationship.

We contribute to the literature as follows. First, this study is one of the first attempts to develop a bloat measure from CD&A sections using generative AI. While Kim et al. (2024) are the first to develop the concept of

bloat and apply it to a different set of disclosures, we document the explanatory power of bloat in an executive compensation context. A key advantage of LLMs is that they ‘understand’ the textual content by contextualizing it against the sections in which that text is embedded. We further contribute to the literature on natural language processing in accounting by showing that this approach complements existing textual metrics. While traditional measures primarily use dictionary based methods to capture surface-level linguistic features such as readability, textual complexity or speech patterns (Li 2008; Loughran and McDonald 2014; Francis et al. 2020), our measure draws on the attention mechanism inherent to LLMs to capture contextual irrelevance, a dimension that is incrementally informative over existing approaches. Our results have direct implications for regulators, such as the SEC, by highlighting that longer or more complex executive compensation disclosures do not necessarily enhance transparency, and for investors, who can use generative AI models to analyze the most important parts of the compensation disclosure sections faster and more easily.

2. Literature review

2.1. Optimal contracting, rent extraction, and excess executive compensation

A central question in corporate governance research is whether executive pay reflects efficient contracting or managerial rent extraction. The efficient contracting view, rooted in classical agency theory (Jensen and Meckling 1976), posits that boards design compensation packages to align managerial incentives with those of shareholders. From this perspective, pay levels simply reflect a competitive market for talent and economic determinants like firm size and value (Gabaix and Landier 2008).

In contrast, the managerial power or rent extraction view (Bebchuk and Fried 2003) argues that executive compensation is not a solution to the agency problem, but a consequence of it. Specifically, it posits that executives can earn excessive levels of compensation due to their power over the board. In line with this view, previous research has shown that excess compensation is associated with weak governance, as proxied by entrenched boards or CEO-chair duality (Core et al. 1999).

A central challenge to the rent extraction view is explaining how this practice persists in the face of mandatory compensation disclosures to enable investor monitoring. One potential mechanism is obfuscation, i.e., increasing the redundancy, complexity, and length of the provided information to draw attention away from key details and making it difficult for investors with limited attention and cognitive processing constraints to draw accurate conclusions (Blankespoor et al. 2019, 2020). In other words, if investors cannot easily process the disclosed compensation information, they are less likely to thoroughly evaluate undeserved excess pay.

In the United States, investors mainly rely on the CD&A section of the proxy statement to understand executive compensation packages. While mandated by the SEC to improve clarity and transparency (SEC 2011), the CD&A in practice is often long and complex. This raises the question whether CD&A sections contain bloat, i.e., irrelevant, boilerplate, and redundant content, and whether the extent of bloat is associated with excess compensation. Although earlier studies have attempted to identify obfuscation in other disclosure contexts using traditional textual analysis methods, such as readability or sentiment (Li 2008; Laksmana et al. 2012; Lawrence 2013), these approaches are not designed to assess contextual irrelevance.

2.2. Bloat and generative AI

Financial disclosures have grown substantially in length over the past two decades (Lesmy et al. 2019), driven mainly by regulation as well as firms voluntarily adding extra information. While increased disclosure should theoretically enhance transparency and reduce information asymmetry (Leuz and Verrecchia 2000), in practice, excessive length may reduce the informativeness and usefulness of financial reporting by overwhelming readers with irrelevant information. This phenomenon, where disclosures become unnecessarily wordy and include redundant information, is referred to as bloat (Kim et al. 2024). Bloat is especially problematic given that individuals have limited cognitive resources to process information (Albuquerque et al. 2024) and investors have restricted attention to analyzing complete financial data (Kim et al. 2024), meaning that they cannot evaluate every detail, particularly when disclosures are repetitive or strategically ambiguous.

Traditional textual analysis methods, such as document readability and length (e.g., Loughran and McDonald 2014), sentiment analysis (e.g., De Amicis et al. (2021)), and topic modelling (e.g., Dyer et al. (2017)), provide valuable insights into corporate reports. However, these tools use various dictionaries and rule-based approaches to determine the overall sentiment of the text without considering the relationship between words and sentences (Leippold 2023). As a result, they capture different aspects of textual complexity but are not designed to assess contextual relevance.

The introduction of generative AI represents a methodological shift and overcomes existing limitations through contextual understanding and reasoning. Unlike traditional NLP models, LLMs are pre-trained on a large set of data, meaning that the model learns patterns in language, such as grammar, word associations, sentence structures, and facts (Floridi and Chiriatti 2020). By filtering out redundancy and producing concise yet fully informative summaries, generative AI models reduce the cost, time, and computing resources required for textual analysis (Huang et al. 2023) and outperform other NLP models in various domains, including question-answering and translation (Brown et al. 2020).

Hence, LLMs' main advantage over traditional models is their ability to perform tasks that involve human-like judgment by understanding the context surrounding each word and the relationships between sentences. These advantages of LLMs are expected to offer a new way of quantifying disclosure bloat in executive compensation disclosures and generate summaries that contain incrementally informative content associated with excess compensation beyond traditional textual measures.

2.3. Bloat and excess executive compensation

Although the structure of performance-based incentives embedded within executive compensation contracts often serves as a mechanism to align managerial actions with shareholder interests (Bebchuk and Fried 2003), if executives benefit from rent extraction, they may have incentives to use disclosure bloat to mask excessive pay. Under the managerial power view, executives may use bloat in the CD&A section to earn compensation that is higher-than-justified by firm performance. First, when financial documents are less readable, shareholders that are limited in time and processing capacity may be less willing to thoroughly evaluate and scrutinize the compensation details, even if the information is relevant (Bloomfield 2002; Hooghiemstra et al. 2017). Second, bloat may focus investors' attention on qualitative narratives framed by the firm, rather than relevant hard information. Consequently, we hypothesize:

H1: Disclosure bloat in the CD&A sections of proxy statements is positively associated with excess executive compensation.

This hypothesis is not without tension, because not all bloat is necessarily opportunistic, as compliance with regulatory requirements and legal risk management may require firms to include lengthy and redundant text to their disclosures.

2.4. Governance mechanisms as bloat determinants

While high executive compensation may be concealed by bloat, executives' power over their compensation may be influenced by the strength of the corporate governance systems in place. Although monitoring and higher shareholder power can limit executives' opportunistic behaviour (Eisenhardt 1989; Ruiz-Verdú 2008), CEO duality or weak board oversight can facilitate obfuscation (Core et al. 1999; Basu et al. 2007), in line with managerial power theory. Moreover, directors may support executive-friendly pay because they value their board positions, are influenced by personal relationships and loyalty to CEOs, and benefit directly from the rewards CEOs can offer (Van Essen et al. 2015). Consequently, a less powerful and independent board may allow executives to exert greater influence over the compensation disclosure content and increase its length, which leads to higher bloat.

Thus, governance mechanisms can be possible determinants of disclosure bloat. Particularly, strong governance mechanisms, such as board gender diversity, should encourage more concise and transparent reporting, hence reducing bloat, while weaker governance, such as CEOs holding both CEO and board chair roles, can increase the opportunity for obfuscation and bloat. These arguments lead to the second hypothesis:

H2a: Board gender diversity is negatively associated with disclosure bloat.

H2b: CEO-Chair duality is positively associated with disclosure bloat.

3. Data and methodology

3.1. Data source and sample

The dataset comprises CD&A sections extracted from DEF 14A proxy filings in the SEC's EDGAR database for S&P 1500 firms between 2011 and 2018, 7,786 firm-year observations. Data regarding the executive compensation and control variables are obtained from Compustat, ExecuComp, and BoardEx. Missing values are further retrieved from Refinitiv Eikon and the company's annual reports.

3.2. Variable measurement and summaries

Variable measurement consists of three main sets of variables and the procedure to construct the summaries. All continuous variables are winsorized at the 1st and 99th percentiles to mitigate the influence of outliers, except when bounded between 0 and 1.

3.2.1. Measurement for excess compensation

Excess executive compensation is measured using the residual of actual compensation compared to the predicted compensation based on the economic determinants, following prior literature by Core et al. (2008) using two different compensation measures. In a robustness check, excess executive compensation is estimated using the Larcker et al. (2011) methodology.

3.2.2. Generating the summaries

Recent advances in machine learning have improved natural language processing, with Transformer-based models becoming some of the most successful architectures to date. The Transformer is a type of deep learning model that uses a neural network architecture, processing sequences through multi-head self-attention to capture relationships between all parts of the input simultaneously (Vaswani et al. 2017). Similar to human thinking, the Transformer model, unlike previous deep learning models, pays attention to the most relevant words when understanding a

sentence. For example, when reading 'Amsterdam is the capital of __,' the model predicts 'The Netherlands' by recognizing the relationship between 'Amsterdam' and 'capital'.

LLMs generate summaries by rephrasing ideas based on context and instruction (prompt), not by copying text. The model writes a shorter version using new sentences, and because the model is trained on vast amounts of data, including financial documents, it can infer what is relevant to investors, even when the language is complex. Therefore, LLMs can tailor summaries for investors and help them understand the CD&A sections more quickly and easily.

Summaries are generated using the DeepSeek-r1-distill-qwen-7b model hosted on a local machine by dividing each CD&A section into 30,000-character chunks and using the following prompt: *Write an investor-friendly summary of all relevant information in this document. The summary should be highly informative and detailed. Do not skip important data or context – include all major figures, policies, and justifications.*

The process works as follows: the CD&A files are split into chunks on a local machine. These chunks, accompanied by a standardized prompt and the model's role, are then sent individually to the LLM. The communication and data exchange between the local and hosted machines occur using Python and an Application Programming Interface (API), which is a set of rules and protocols that allows different applications to interact in a well-documented way (Chen 2025). The LLM processes the inputs remotely and sends the generated summaries back to the local machine, where the outputs are merged and stored.

3.2.3. Measurement of bloat and traditional textual methods

Bloat is measured following the methodology by Kim et al. (2024) as the difference between original text and summary lengths scaled by original text length. A high value of *bloat* implies more redundant and irrelevant information. To assess whether the bloat measure provides incremental information beyond established linguistic features, several traditional textual controls are included: *Fog Index* (Gunning 1952; Li 2008; Kim et al. 2024), *File Length* (Li 2008; Lawrence 2013), *SentimentLM* analysis (Loughran and McDonald 2014; Huang et al. 2014), *Redundancy* (Dyer et al. 2017; Kim et al. 2024), *Boilerplate* language (Dyer et al. 2017; Kim et al. 2024).

3.2.4. Control variables

To control for the variables that may impact bloat and influence the relationship between the variables of interest, building on prior literature, several controls are incorporated: firm size (Li 2008; Carter et al. 2016), leverage (Carter et al. 2016), file size, industry, firm, and year-fixed effects. When the models are estimated using industry-fixed effects instead of firm-fixed effects,

additional variables are included as controls: sales growth (Kramer and Matějka 2024), R&D expenses (Albuquerque et al. 2024), and ROE (Carter et al. 2016) to control for cross-industry variation in performance, complexity, and growth that may influence disclosure practices and executive compensation.

3.3. Research design

To answer the research question and test H1, whether bloat is positively associated with excess compensation, the following OLS panel regression with fixed effects is employed:

$$Excess(Comp)_{it} = \beta_0 + \beta_1 Bloat_{it} + \beta_2 TraditionalMethods + \beta_3 FirmControls_{it} + \alpha_i + \delta_t + \theta_i + \varepsilon_{it} \quad (1)$$

where α_i is firm fixed effect, δ_t is year fixed effect, θ_i is industry fixed effects using the Fama and French (1997) 48–industry classification, and ε_{it} is the error term.

To address H2a and H2b and find possible determinants of bloat, board gender diversity and CEO-chair duality are regressed with CEO pay slice (Bebchuk et al. 2011) and board size (Yermack 1996) as additional possible determinants and growth options, volatility, and loss (Kim et al. 2024) as control variables in the OLS regression:

$$Bloat_{it} = \beta_0 + \beta_1 GovernanceMechanisms_{it} + \beta_2 OtherFactors_{it} + \beta_3 FirmFactors_{it} + \delta_t + \theta_i + \varepsilon_{it} \quad (2)$$

Table 1. Descriptive statistics.

Variable	Number of Observations	Mean	Standard Deviation	Q1	Median	Q3
<i>ExcessComp1</i>	6,668	0.000	0.662	−0.297	0.066	0.381
<i>ExcessComp2</i>	6,668	0.000	0.790	−0.428	0.029	0.464
<i>ExcessPay</i>	7,786	−0.942	1.001	−0.362	0.000	0.315
<i>Bloat</i>	7,786	0.874	0.033	0.862	0.888	0.893
<i>Fog</i>	7,786	22.100	1.760	20.900	22.000	23.100
<i>File_Length</i>	7,786	14,017	6,882	9,356	12,998	17,628
<i>File_Size_kb</i>	7,786	107.699	49.903	72.987	101.836	134.690
<i>SentimentLM</i>	7,786	−0.000	0.004	−0.003	0.000	0.003
<i>Redundancy</i>	7,786	0.189	0.058	0.150	0.184	0.224
<i>Boilerplate</i>	7,786	0.016	0.005	0.013	0.015	0.018
<i>Leverage</i>	7,786	0.200	0.195	0.033	0.153	0.306
<i>Firm_Size</i>	7,786	7.651	1.637	6.538	7.590	8.700
<i>Sales_Growth</i>	7,786	0.089	0.229	−0.009	0.062	0.147
<i>R&D_Assets</i>	7,786	0.038	0.078	0.000	0.000	0.039
<i>ROE</i>	7,786	0.075	0.497	0.032	0.108	0.187
<i>BoardDiversity</i>	7,786	0.099	0.136	0.000	0.000	0.200
<i>CEO_Duality</i>	7,786	0.459	0.498	0.000	0.000	1.000
<i>CEO_Pay_Slice</i>	7,786	0.397	0.120	0.331	0.404	0.466
<i>Board_Size</i>	7,786	5.698	1.066	5.000	5.000	6.000
<i>Loss</i>	7,786	0.182	0.390	0.000	0.000	0.000
<i>Volatility</i>	7,526	0.335	0.160	0.223	0.301	0.402
<i>Growth_Options</i>	7,786	2.735	2.346	1.361	1.959	3.108

Note. The sample has 7,786 firm-year observations of S&P 1500 companies from 2011–2018 with non-missing data for the main variables of interest, for which it was possible to extract CD&A sections and each firm had at least 2 years of observations. Since excess compensation calculation involves lagged values, the final number of observations used is 6,668. There are 260 missing observations for *Volatility*.

3.4. Descriptive statistics

Table 1 shows summary statistics for the main variables used in the analysis. The interpretation of excess compensation is the log ratio of actual to predicted compensation. Therefore, *ExcessComp1* indicates that the median excess compensation is approximately 6.8 percent higher than predicted. All three proxy variables for excess executive compensation have a high standard deviation, meaning that the actual pay of some executives is more than their predicted pay, while for others, it is below predicted. Mean value of *Bloat* suggests that, on average, 87.4% of the original document content is not retained in the summary, implying a high level of redundancy in compensation disclosures.

4. Results

4.1. Hypothesis 1 – Disclosure bloat and excess executive compensation

The first analysis tests whether AI-generated summaries contain incrementally informative content beyond traditional textual measures in explaining excess compensation. Using a regression model in Equation (1), Table 2 shows that in column (1), a 0.01 point increase in net positive sentiment of *SentimentLM* ($p < 0.001$) is associated with an 8.89% ($= e^{8.518 \times 0.01} - 1$) increase in excess compensation, keeping other variables constant. While this result suggests that companies often use a more positive tone

in their documents as a strategic tool to manage perception and shift readers' attention away from things they want to hide, aligning with prior research (Huang et al. 2014, 2018), positive sentiment could coincide with the underlying justification for higher pay, such as strong performance.

The positive direction of *File_Length* ($p < 0.001$) and *Res_File_Size* ($p < 0.001$) coefficients is not surprising since larger file lengths and sizes suggest higher textual complexity. Every additional 1,000 words is associated with a 0.7% increase in excess compensation. Moreover, the significance of the *Res_File_Size* implies that firms that use more complex formatting and other textual characteristics beyond the word count are also more likely to provide longer and more detailed information, not for the sake of transparency, but to overwhelm the readers and obscure important information (Li 2008; Lawrence 2013). Therefore, less scrutiny can allow excess executive compensation to go unnoticed.

The coefficient of *Leverage* (-0.257 , $p < 0.001$) indicates that higher debt constrains executives' ability to inflate pay because creditors, particularly banks, act as additional monitors when they provide substantial financing. Hence, increased oversight limits executives' ability to earn excess compensation. The *Fog_Index* coefficient is insignificant ($p > 0.1$), meaning that it cannot explain excess compensation. This conclusion opposes the initial expectation based on Li (2008) but shows that traditional readability metrics fail to capture textual complexity relevant to excess compensation, consistent with Loughran and McDonald (2014).

When industry-fixed effects replace firm-fixed effects in Columns (3) and (4), *SentimentLM*, *Leverage*, and *Bloat* lose significance, and the adjusted R^2 decreases from 60.6% to 9.4%, implying that industry fixed effects provide a much weaker control structure than firm fixed effects. Firm fixed effects absorb all stable, unobserved firm-level heterogeneity, whereas industry fixed effects only control for broad sector differences. As a result, switching to industry fixed effects introduces additional noise and reduces explanatory power, making coefficient estimates less precise.

Table 3 displays the results of Equation (1) with *ExcessComp2* as a dependent variable. The findings are consistent with Table 2. *Firm_Size* becomes significant at the 5% level, previously 10%, suggesting larger firms are positively associated with excess executive compensation due to that larger firms have longer and more complex disclosures, which offer executives greater opportunity for obfuscation (Li 2008).

Similarly, Columns (3) and (4) in Table 3 are consistent with the results in Columns (3) and (4) in Table 2, but the *Fog_Index* coefficient becomes significant ($p < 0.001$) when industry-fixed effects are included, which implies that the variation in *Fog_Index* that matters for excess compensation occurs between firms, not within a firm over time. This finding differs from previous outcomes but aligns with Li (2008), who argues that higher textual

Table 2. Test of H1: Excess Compensation Consequences of Disclosure Bloat.

Dependent Variable =	<i>ExcessComp1</i>			
	(1)	(2)	(3)	(4)
<i>Bloat</i>		0.520*** (0.200)		0.152 (0.262)
<i>Fog_Index</i>	-0.000 (0.009)	-0.000 (0.009)	0.009 (0.007)	0.009 (0.007)
<i>Redundancy</i>	0.340 (0.249)	0.326 (0.247)	0.304 (0.372)	0.371 (0.299)
<i>SentimentLM</i>	8.518*** (3.148)	8.285*** (3.160)	-4.184 (3.530)	-4.233 (3.504)
<i>File_Length</i>	0.007*** (0.002)	0.007*** (0.005)	0.016*** (0.004)	0.015*** (0.004)
<i>Boilerplate</i>	0.200 (3.188)	1.047 (3.223)	-1.233 (3.311)	-1.075 (3.304)
<i>Firm_Size</i>	0.074* (0.045)	0.074* (0.045)	0.024 (0.017)	0.024 (0.017)
<i>Leverage</i>	-0.257*** (0.088)	-0.259*** (0.089)	0.181 (0.091)	0.181* (0.091)
<i>Res_File_Size</i>	0.002*** (0.001)	0.002*** (0.001)	0.007*** (0.001)	0.007*** (0.001)
<i>Sales_Growth</i>			0.268*** (0.038)	0.267*** (0.038)
<i>R&D_Assets</i>			0.523 (0.542)	0.523 (0.542)
<i>ROE</i>			0.005 (0.014)	0.005 (0.014)
Firm FE	Yes	Yes	No	No
Year FE	Yes	Yes	Yes	Yes
Industry FE	No	No	Yes	Yes
Adj. R^2	0.606	0.606	0.094	0.094
Within R^2	0.015	0.016	0.074	0.074
N	6,668	6,668	6,668	6,668

Note. Presented are OLS estimates of the main research model described in Equation (1) with excess compensation calculated using *tdc1* as a dependent variable. ***, **, * represent significance at the 0.01, 0.05 and 0.10 levels. Standard errors are in parentheses and clustered at the Fama-French 48-industry level in Columns (1) and (2) and at the firm level in Columns (3) and (4). Year-fixed effects are included in all regressions. Since excess compensation calculation involves lagged values, the final number of observations used is 6,668.

complexity should enable executives to hide details and increase their compensation.

The results align with theoretical expectations. As argued before, most traditional measures rely on keywords and rules, while ignoring contextual meaning (Leippold 2023). Since compensation contracts are often written in verbose and complex language (Albuquerque et al. 2024), in line with the management obfuscation hypothesis, shareholders are less likely to thoroughly evaluate compensation structures due to the time constraints they face, leading to excess compensation. In contrast, LLMs are designed to grasp the contextual significance, as Floridi and Chiriatti (2020) argue, allowing them to filter out unnecessary and repetitive language and extract essential informational content from complex disclosures.

Table 3. Test of H1: Excess compensation consequences of disclosure bloat.

Dependent Variable =	<i>ExcessComp2</i>			
	(1)	(2)	(3)	(4)
<i>Bloat</i>		0.944*** (0.278)		0.317 (0.370)
<i>Fog_Index</i>	0.009 (0.012)	0.010 (0.012)	0.022** (0.008)	0.022** (0.008)
<i>Redundancy</i>	-0.040 (0.292)	-0.066 (0.289)	0.231 (0.497)	0.222 (0.494)
<i>SentimentLM</i>	10.302** (4.025)	9.880** (4.030)	-3.990 (3.909)	-4.093 (3.822)
<i>File_Length</i>	0.010*** (0.002)	0.009*** (0.002)	0.015*** (0.004)	0.014*** (0.004)
<i>Boilerplate</i>	-1.468 (3.822)	0.689 (3.852)	-1.868 (3.473)	-1.538 (3.427)
<i>Firm_Size</i>	0.090** (0.040)	0.090** (0.040)	0.023 (0.017)	0.023 (0.017)
<i>Leverage</i>	-0.218** (0.104)	-0.222** (0.104)	0.186* (0.104)	0.186* (0.104)
<i>Res_File_Size</i>	0.002*** (0.001)	0.003*** (0.001)	0.007*** (0.001)	0.007*** (0.001)
<i>Sales_Growth</i>			0.333*** (0.039)	0.333*** (0.039)
<i>R&D_Assets</i>			0.200 (0.436)	0.200 (0.436)
<i>ROE</i>			0.015 (0.020)	0.015 (0.020)
Firm FE	Yes	Yes	No	No
Year FE	Yes	Yes	Yes	Yes
Industry FE	No	No	Yes	Yes
Adj. R ²	0.498	0.499	0.069	0.069
Within R ²	0.012	0.014	0.056	0.056
N	6,668	6,668	6,668	6,668

Note. Presented are OLS estimates of the main research model described in Equation (1) with excess compensation calculated using *tdc2* as a dependent variable. ***, **, * represent significance at the 0.01, 0.05 and 0.10 levels. Standard errors are in parentheses and clustered at the Fama-French 48–industry level in Columns (1) and (2) and at the firm level in Columns (3) and (4). Year-fixed effects are included in all regressions. Since excess compensation calculation involves lagged values, the final number of observations used is 6,668.

Following this rationale, the added explanatory power of the *Bloat* variable is a reasonable outcome of using a technology that more accurately replicates human judgment and intent when determining relevant content (Dong et al. 2024). Hence, the main implication of H1 is that LLMs can improve the quality of corporate disclosures and help address the principal-agency problem by assisting stakeholders in better understanding the company's financial statements and reports.

To address whether summaries predict excess compensation, given that *Bloat* varies over time within firms, I use column (2) in Tables 2, 3 with firm-fixed effects to test H1. The coefficient of *Bloat* in column (2) of Table 2 is positive and significant at a 1% significance level; a 0.01 unit increase in *Bloat* is associated with a 0.52% ($= e^{0.520 \times 0.01} - 1$) increase in *ExcessComp1* and a 0.95%

($= e^{0.944 \times 0.01} - 1$) increase in *ExcessComp2*, when all other variables are kept constant. Hence, firms with more bloated CD&A disclosures are more likely to give executives compensation higher than what is justified by the firm's performance. As a result, H1 is not rejected.

The result is consistent with expectations from prior literature. According to agency theory, executives are incentivized to exploit their managerial power and camouflage the structure of their compensation (Bebchuk and Fried 2003), and bloated disclosures can be a mechanism to earn excess pay without detection by investors. Since investors and shareholders face cognitive and informational constraints (Albuquerque et al. 2024; Kim et al. 2024), they may give less scrutiny to compensation disclosures. Laksmana et al. (2012) also find lower CD&A readability in the 2007 proxy statements when the CEO's compensation exceeds what is justified by the firm's performance.

Similar to the results in Kim et al. (2024), *Redundancy* and *Boilerplate* are not significant in any of the columns. This finding further highlights the advantage of generative AI in better understanding the textual complexities. Nevertheless, the size and significance of the *SentimentLM*, *File_Length*, and *Res_File_Size* variables stay mostly constant, and the within R² does not substantially increase after incorporating the *Bloat* variable into the regression, even though the coefficient of *Bloat* is statistically significant. All of these suggest that the *Bloat* does not correct for omitted variable bias; instead, it reflects a unique but very small component of disclosure complexity that is not fully explained by sentiment or document length.

Untabulated tests on sentiment tone reveal that the negative sentiment coefficient is insignificant ($p > 0.1$), while the positive sentiment coefficient is significant ($p < 0.001$) in the main regression results. This suggests that firms use a more positive tone to obscure information since highlighting positive events can make investors judge higher executive compensation less, consistent with Huang et al. (2018), who show that managers use a positive tone to enhance earnings management.

The robustness check, where excess executive compensation is calculated using an alternative proxy of the Larcker et al. (2011) model, shows that the estimate of *Bloat* remains significant and its magnitude is similar to the one in Column (2) of Table 2. Moreover, as before, *Bloat* is significant only with firm-fixed effects, not with industry-fixed effects.

Regarding the economic magnitude of the *Bloat* variable, using the estimates from column (2) of Table 3, a one standard deviation increase in *Bloat* is associated with a 3.16% ($= e^{0.033 \times 0.944} - 1$) increase in excess executive compensation. Although 3.16% might appear small, a slight change in executive pay can amount to thousands of dollars for CEOs in larger firms, which highlights the practical importance and LLM's potential for increasing transparency and accountability in corporate governance.

4.2. Hypothesis H2a and H2b – Determinants of bloat

We find no significant relationship between *BoardDiversity* or *CEO_Duality* and *Bloat*, suggesting that formal governance structures may not effectively constrain redundancy. Hence, H2a and H2b are rejected. Overall, the model explains only 0.9% of the variation in *Bloat*, further indicating the heterogeneous complexity of disclosure practices. However, the findings align with managerial power theory (Eisenhardt 1989) that since board members value their positions and personal relationships with CEOs, they are less likely to disagree with the CEO's pay structure and reduce redundant compensation content.

5. Conclusion and discussion

This research explored whether large language models can identify redundant content in the CD&A sections of proxy statements and whether this redundancy is associated with excess executive compensation. The study introduces a new way of understanding excess executive compensation using a quantifiable, verbose content measure – *bloat* – beyond traditional textual metrics. To address the research question of how a measure of compensation disclosure bloat can be developed using generative AI and its relationship with excess executive compensation, fixed-effects panel regression models were created for 7,786 firm-year observations covering S&P 1500 companies between 2011–2018.

The results show that the bloat measure significantly explains variation in excess executive compensation, and it captures the part of excess compensation that traditional natural language processing models cannot explain. The regression results from Equation (1) are consistent with H1, which suggests that bloated CD&A sections are associated with higher excess executive compensation. The rationale is that according to agency theory, executives' goals are to maximize personal benefits, and bloat allows them to obscure key pay details by overwhelming investors with irrelevant information and thus justifying excess compensation because investors have limited cognitive resources and restricted attention to analyzing full financial data (Albuquerque et al. 2024; Kim et al. 2024).

Interestingly, including the industry-fixed effects rather than the firm-fixed effects turned a significant bloat coefficient into an insignificant one. This implies that the effect of the bloat variable is driven by within-firm variation over time, and a large portion of the variation in excess compensation is explained by firm-specific factors that do not vary much over time.

Contrary to H2a and H2b expectations, we did not find strong evidence that governance mechanisms are associated with less bloat. This may be due to measurement limitations or the choice of governance proxies. In addition, compensation disclosure practices may be shaped more by institutional or advisory norms than by

formal governance structures, which can be a focus for future studies.

5.1. Implications

This research improves both social and scientific understanding of corporate disclosure by demonstrating the value of LLMs in identifying complex textual relationships that traditional metrics often fail to identify (Leipold 2023). It is the first study to use LLMs to quantify disclosure bloat in corporate governance and examine its relationship with excess executive compensation. The results can offer a solution to a longstanding problem of excessive executive compensation by highlighting the potential of generative AI to address it.

Our findings suggest important implications for regulators by showing that lengthy or more complex executive compensation disclosures do not automatically improve transparency. They also matter for investors, who can leverage generative AI to analyze the key elements of the compensation disclosure sections faster and more easily. Finally, corporate boards should emphasize the significance of clear and less complex communication in financial documents.

5.2. Limitations and future research

We acknowledge some limitations that offer opportunities for future research. First, the output of the LLM varies depending on the model and prompt used. As a result, the length of the summarized documents reflects the LLM's judgment of contextual relevance and can change depending on the assigned task to a specific LLM, which may also affect the measure of the bloat variable. Future research can utilize different and more advanced models to assess whether the findings are generalizable.

Moreover, the analysis is restricted to S&P 1500 firms from 2011–2018, limiting generalizability to smaller or international firms with different disclosure practices. Expanding the sample and improving CD&A data extraction would strengthen external validity. Future research can also examine non-financial reports, such as sustainability reports, to investigate whether management obscures relevant information when the content is less numerical.

Although our empirical analysis is based on U.S. firms, the main message that quality matters more than length is highly relevant to the European context. Future research could apply the bloat measure to European remuneration reports to see whether the results hold under the Shareholder Rights Directive.

Finally, as with most empirical corporate governance research, absent an exogenous shock endogeneity remains a concern. We therefore do not claim that our associations imply causality. Nevertheless, the findings provide an important starting point for deepening the understanding of how the complexity of text in executive pay disclosures may relate to governance outcomes and remain valuable for theory development and policy discussions.

- **L. Burduli – Lizi Burduli** holds a MScBA in Accounting & Financial Management from RSM, Erasmus University and is currently an auditor at Deloitte.
- **Dr. S. Kramer – Stephan Kramer** is a Professor of Financial Decision Making and Control at RSM, Erasmus University.

This article is based on Lizi Burduli's master's thesis. This makes her one of the winners of the MAB Thesis Award 2025.

References

- Albuquerque A, Carter ME, Guo ZM, Lynch LJ (2024) Complexity of CEO compensation packages. *Journal of Accounting and Economics* 79(1): 101709. <https://doi.org/10.1016/j.jacceco.2024.101709>
- Basu S, Hwang LS, Mitsudome T, Weintrop J (2007) Corporate governance, top executive compensation and firm performance in Japan. *Pacific-Basin Finance Journal* 15(1): 56–79. <https://doi.org/10.1016/j.pacfin.2006.05.002>
- Bebchuk LA, Fried JM (2003) Executive compensation as an agency problem. *Journal of Economic Perspectives* 17(3): 71–92. <https://doi.org/10.1257/089533003769204362>
- Bebchuk LA, Cremers KM, Peyer UC (2011) The CEO pay slice. *Journal of Financial Economics* 102(1): 199–221. <https://doi.org/10.1016/j.jfineco.2011.05.006>
- Blankespoor E, DeHaan E, Marinovic I (2020) Disclosure processing costs, investors' information choice, and equity market outcomes: A review. *Journal of Accounting and Economics* 70(2–3): 101344. <https://doi.org/10.1016/j.jacceco.2020.101344>
- Blankespoor E, DeHaan ED, Wertz J, Zhu C (2019) Why do individual investors disregard accounting information? The roles of information awareness and acquisition costs. *Journal of Accounting Research* 57(1): 53–84. <https://doi.org/10.1111/1475-679X.12248>
- Bloomfield RJ (2002) The 'incomplete revelation hypothesis' and financial reporting. *Accounting Horizons* 16: 233–243. <https://doi.org/10.2308/acch.2002.16.3.233>
- Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler DM, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I, Amodei D (2020) Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33: 1877–1901. <https://doi.org/10.48550/arXiv.2005.14165>
- Carter ME, Li L, Marcus AJ, Tehranian H (2016) Excess pay and deficient performance. *Review of Financial Economics* 30: 1–10. <https://doi.org/10.1016/j.rfe.2015.08.003>
- Chen M (2025) [February 24] What is an API (Application Programming Interface)? <https://www.oracle.com/nl/cloud/cloud-native/api-management/what-is-api>
- Core JE, Holthausen RW, Larcker DF (1999) Corporate governance, chief executive officer compensation, and firm performance. *Journal of Financial Economics* 51(3): 371–406. [https://doi.org/10.1016/S0304-405X\(98\)00058-0](https://doi.org/10.1016/S0304-405X(98)00058-0)
- Core J, Guay W, Larcker D (2008) The power of the pen and executive compensation. *Journal of Financial Economics* 88: 1–25. <https://doi.org/10.1016/j.jfineco.2007.05.001>
- De Amicis C, Falconieri S, Tastan M (2021) Sentiment analysis and gender differences in earnings conference calls. *Journal of Corporate Finance* 71: 101809. <https://doi.org/10.1016/j.jcorpfin.2020.101809>
- Directive 2007/36/EC (2007) Directive 2007/36/EC of the European Parliament and of the Council of 11 July 2007 on the exercise of certain rights of shareholders in listed companies. <https://eur-lex.europa.eu/eli/dir/2007/36/oj/eng>
- Dong MM, Stratopoulos TC, Wang VX (2024) A scoping review of ChatGPT research in accounting and finance. *International Journal of Accounting Information Systems* 55: 100715. <https://doi.org/10.1016/j.accinf.2024.100715>
- Dyer T, Lang M, Stice-Lawrence L (2017) The evolution of 10-K textual disclosure: Evidence from Latent Dirichlet Allocation. *Journal of Accounting and Economics* 64(2–3): 221–245. <https://doi.org/10.1016/j.jacceco.2017.07.002>
- Eisenhardt KM (1989) Agency theory: An assessment and review. *Academy of Management Review* 14(1): 57–74. <https://doi.org/10.5465/amr.1989.4279003>
- Fama EF, French KR (1997) Industry costs of equity. *Journal of Financial Economics* 43(2): 153–193. [https://doi.org/10.1016/S0304-405X\(96\)00896-3](https://doi.org/10.1016/S0304-405X(96)00896-3)
- Floridi L, Chiriatti M (2020) GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines* 30: 681–694. <https://doi.org/10.1007/s11023-020-09548-1>
- Francis BB, Shohfi T, Xin D (2020) Gender and earnings conference calls. SSRN 25: 2020. <https://doi.org/10.2139/SSRN.3473266>
- Gabaix X, Landier A (2008) Why has CEO pay increased so much? *The Quarterly Journal of Economics* 123(1): 49–100. <https://doi.org/10.1162/qjec.2008.123.1.49>
- Gunning R (1952) *The technique of clear writing*. McGraw-Hill.
- Hooghiemstra R, Kuang YF, Qin B (2017) Does obfuscating excessive CEO pay work? The influence of remuneration report readability on say-on-pay votes. *Accounting and Business Research* 47(6): 695–729. <https://doi.org/10.1080/00014788.2017.1300516>
- Huang X, Teoh SH, Zhang Y (2014) Tone management. *The Accounting Review* 89(3): 1083–1113. <https://doi.org/10.2308/accr-50684>
- Huang X, Krishnan S, Lin P (2018) Tone analysis and earnings management. *Journal of Accounting and Finance* 18(8): 46–61. <https://doi.org/10.33423/jaf.v18i8.110>
- Huang AH, Wang H, Yang Y (2023) FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research* 40(2): 806–841. <https://doi.org/10.1111/1911-3846.12832>

- Jensen MC, Meckling WH (1976) Theory of the firm: Managerial behavior, agency costs, and ownership structure. *Journal of Financial Economics* 3(4): 305–360. [https://doi.org/10.1016/0304-405X\(76\)90026-X](https://doi.org/10.1016/0304-405X(76)90026-X)
- Kim A, Muhn M, Nikolaev VV (2024) Bloated disclosures: can ChatGPT help investors process information? *Chicago Booth Research Paper* (23-07): 2023–59. <https://doi.org/10.48550/arXiv.2306.10224>
- Kramer S, Matějka M (2024) Disturbing the Quiet Life? Competition and CEO Incentives. *The Accounting Review* 99(2): 279–305. <https://doi.org/10.2308/TAR-2022-0393>
- Laksmana I, Tietz W, Yang YW (2012) Compensation discussion and analysis (CD&A): Readability and management obfuscation. *Journal of Accounting and Public Policy* 31(2): 185–203. <https://doi.org/10.1016/j.jaccpubpol.2011.08.003>
- Larcker DF, Ormazabal G, Taylor DJ (2011) The market reaction to corporate governance regulation. *Journal of Financial Economics* 101(2): 431–448. <https://doi.org/10.1016/j.jfineco.2011.03.002>
- Lawrence A (2013) Individual investors and financial disclosure. *Journal of Accounting and Economics* 56(1): 130–147. <https://doi.org/10.1016/j.jacceco.2013.05.001>
- Leippold M (2023) Sentiment spin: Attacking financial sentiment with GPT-3. *Finance Research Letters* 55: 103957. <https://doi.org/10.1016/j.frl.2023.103957>
- Lesmy D, Muchnik L, Mugerma Y (2019) Doyoureadme? temporal trends in the language complexity of financial reporting. *Temporal Trends in the Language Complexity of Financial Reporting* (September 26, 2019). <https://doi.org/10.2139/ssrn.3469073>
- Leuz C, Verrecchia RE (2000) The economic consequences of increased disclosure. *Journal of Accounting Research*: 91–124. <https://doi.org/10.2307/2672910>
- Li F (2008) Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics* 45(2–3): 221–247. <https://doi.org/10.1016/j.jacceco.2008.02.003>
- Loughran T, McDonald B (2014) Measuring readability in financial disclosures. *The Journal of Finance* 69(4): 1643–1671. <https://doi.org/10.1111/jofi.12162>
- Ruiz-Verdú P (2008) Corporate governance when managers set their own pay. *European Financial Management* 14(5): 921–943. <https://doi.org/10.1111/j.1468-036X.2008.00465.x>
- SEC [U.S. Securities and Exchange Commission] (2011) Executive compensation. <https://www.sec.gov/answers/excomp.htm>
- Van Essen M, Otten J, Carberry EJ (2015) Assessing managerial power theory: A meta-analytic approach to understanding the determinants of CEO compensation. *Journal of Management* 41(1): 164–202. <https://doi.org/10.1177/0149206311429378>
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. *Advances in Neural Information Processing Systems* 30. <https://doi.org/10.48550/arXiv.1706.03762>
- Yermack D (1996) Higher market valuation of companies with a small board of directors. *Journal of Financial Economics* 40(2): 185–211. [https://doi.org/10.1016/0304-405X\(95\)00844-5](https://doi.org/10.1016/0304-405X(95)00844-5)