Behavioral Finance Analytics with Large Language Models: Exploring the Relationship Between Investor Self-Perception, Decision Biases, and Market Outcomes

Paul M. Carr
Department of Computer Science, University of New Hampshire, Durham, NH, USA.
paul2004@unh.edu

Abstract

The integration of large language models into behavioral finance opens a novel frontier for studying how investors' self-perception interacts with cognitive biases to shape aggregate market outcomes. Traditional behavioral research relies on surveys and experimental proxies that lack ecological richness and scalability. This paper presents a system-level examination of behavioral finance analytics powered by large language models, focusing on the extraction of self-perception narratives, bias signatures, and their empirical linkages to trading volumes, volatility, and bubble formation. We articulate a multi-layered architecture that ingests heterogeneous textual streams—social media discussions, earnings call transcripts, and financial news—and deploys fine-tuned transformer models to detect overconfidence, loss aversion, herding, and confirmation biases embedded in first-person investor discourse. The analysis highlights structural trade-offs among model capacity, inference cost, and analytical fidelity, as well as the deployment challenges of streaming pipelines and model drift. Governance dimensions receive sustained attention, including fairness disparities arising from demographic imbalances in training corpora, robustness against adversarial narrative manipulation, and interpretability requirements mandated by evolving AI regulatory frameworks. We further discuss sustainability considerations such as energy footprints of large foundation models versus distilled variants. The paper concludes with policy implications for market surveillance, systemic risk detection, and the design of transparent AI-assisted behavioral analytics. By linking fine-grained linguistic self-representation to macro-financial dynamics, the proposed analytics framework opens pathways toward more resilient, fair, and context-aware financial decision-support systems.

Keywords

behavioral finance, large language models, investor self-perception, cognitive biases, market outcomes, sentiment analysis, AI governance, fairness.

1. Introduction

The behavioral finance literature has long established that investors systematically deviate from the rationality assumptions of classical finance, driven by cognitive biases, emotional influences, and social dynamics. Foundational work on prospect theory documented how individuals evaluate gains and losses asymmetrically, fundamentally challenging expected utility maximization [1]. Subsequent research demonstrated that aggregate investor sentiment, defined as beliefs about future cash flows and risks not warranted by fundamentals, exerts measurable effects on asset prices, leading to mispricing and prolonged deviations from intrinsic values [2]. A comprehensive treatment of how psychological heuristics create market inefficiencies across a wide range of settings reinforces the need to move beyond price-based measures and examine the cognitive processes that generate trading decisions [3]. Parallel to these theoretical advances, the last decade has witnessed explosive growth in the volume and variety of textual data produced by market participants, ranging from social media posts and investment forum discussions to corporate disclosures and analyst reports. The availability of such data intersects with the remarkable capabilities of large language models (LLMs), exemplified by architectures such as BERT, which have revolutionized natural language understanding by capturing nuanced semantic and contextual signals [4]. Despite the rich potential, the systematic integration of LLM-driven analytics into behavioral finance remains nascent, particularly regarding the granular measurement of investor self-perception—how investors view their own competence, risk tolerance, and decision-making style—and the tracing of those self-perceptions to systematic biases and eventual market outcomes.

Contemporary behavioral finance analytics must confront a multi-scale challenge: linking individual-level cognitive architectures, manifested through language, to emergent market phenomena such as excess volatility, bubbles, and crashes. While earlier sentiment analysis approaches used dictionary-based methods to quantify optimism or pessimism, they cannot capture the layered rhetorical strategies through which investors construct self-narratives, rationalize past decisions, or project overconfidence. LLMs, with their capacity for in-context learning and fine-tuning on domain-specific corpora, offer a pathway to decode these narratives at scale. However, the deployment of such analytics is not merely a computational problem; it raises profound system-level questions about architecture design, data governance, fairness across heterogeneous investor populations, robustness to adversarial content, interpretability, and the sustainability of large-scale model serving. This paper presents a holistic, system-oriented examination of behavioral finance analytics with LLMs, foregrounding the relationship between investor self-perception, decision biases, and market outcomes while attending to the structural trade-offs, governance frameworks, and policy implications that accompany such sociotechnical infrastructures.

2. Foundations of Behavioral Finance and Textual Analysis

To situate the LLM-powered analytics framework, it is necessary to review how behavioral finance has conceptualized decision biases and how textual analysis has been adopted to measure psychological constructs. Early behavioral models identified overconfidence as a pervasive bias: investors overestimate the precision of their private information and trade excessively, generating elevated volume and volatility [7]. A related literature formalized how biased self-attribution leads to momentum and long-run reversals as investors overreact to confirming signals and underweight contradictory evidence [8]. While these studies primarily relied on trade and return data, they established the theoretical imperative to measure belief distortions directly. Concurrently, a stream of research began to exploit large-scale text corpora to quantify the information environment. Pioneering work on financial text analysis demonstrated that carefully constructed dictionaries tailored to business language outperform generic sentiment lexicons, revealing meaningful relationships between the tone of regulatory filings and subsequent returns [5]. The introduction of internet stock message boards provided a new lens, showing that disaggregated bullish and bearish message volumes contain information relevant to volatility and trading activity, even after controlling for traditional news [9]. Further studies solidified the connection between social media mood extracted from platforms such as Twitter and daily stock market movements, indicating that collective emotional states can influence aggregate financial dynamics [10]. The "wisdom of crowds" effect was also documented, with aggregated stock opinions from social media platforms predicting earnings surprises and returns [11]. These empirical regularities underscore the value of unstructured textual data for behavioral finance.

Nevertheless, the dominant textual analysis tradition has focused on coarse sentiment polarity rather than on the nuanced self-referential language that reveals investors' cognitive biases. Forward-looking statements in corporate filings were analyzed with naive Bayesian methods to extract managerial tone, linking linguistic optimism to future earnings and investment [12]. Content analysis of media language further showed that pessimistic words in news columns predict downward pressure on stock prices, and that the linguistic tone of news captures a dimension of investor sentiment distinct from simple word counts [13]. Methodological refinements illustrated that the power of individual words, weighted by their importance in context, can sharpen the measurement of financial text [14]. Analyst report language was shown to contain information beyond quantitative forecasts, with the tone of textual discussions predicting future earnings and recommendation changes [15]. These advances collectively demonstrate that financial text is a rich, high-dimensional signal. Yet a critical gap remains: the explicit modeling of self-perception—the first-person narratives through which investors characterize their own abilities and strategies—has been insufficiently examined. Research in human-computer interaction has explored how individuals construct and express self-identity in digital environments, highlighting the importance of self-reflective language as a window into cognitive and motivational states [6]. This insight, when transplanted into financial contexts, suggests that investors' self-descriptions contain signals of overconfidence, regret avoidance, and attribution biases that go beyond standard sentiment measures. Capturing such signals requires architectures that can parse not just what is said about assets, but how investors position themselves as capable, cautious, or trend-following agents.

3. System Architecture for LLM-Based Behavioral Analytics

Designing an analytics system that transforms heterogeneous financial text streams into validated bias indicators and market outcome predictions demands a multi-layer architecture with careful attention to data ingestion, model orchestration, and downstream inference. The architecture begins with a data acquisition layer capable of ingesting streaming data from diverse sources: real-time Reddit forums, Twitter feeds, earnings call transcripts, broker reports, and investor surveys. Because financial text is domain-specific, preprocessing must handle jargon, abbreviations, numerical entities, and sarcasm, while preserving the first-person constructions that encode self-perception. The ingestion pipeline must also enforce lineage tracking and consent management, given that many sources involve individual investor communications subject to privacy expectations. Following preprocessing, the system routes text through a series of specialized LLM components. The initial stage may employ a lightweight, pre-trained encoder model to filter irrelevant content, reducing the downstream computational burden. A subsequent fine-tuned transformer model, such as a domain-adapted variant of BERT, is tasked with extracting self-perception markers: phrases indicating high or low self-confidence, comparative self-evaluations, and attributions of success or failure. This requires training on carefully annotated corpora where human raters label spans of text according to psychological dimensions derived from behavioral finance theory.

The core analytics engine integrates multiple LLM capabilities. Zero-shot and few-shot prompting with large autoregressive models such as GPT-3 can be used to generate structured summaries of investor self-narratives, identifying patterns of overconfidence, loss aversion, or herding orientation without extensive retraining [17]. Foundation models, which capture broad linguistic competencies, can be adapted to financial domains through parameter-efficient fine-tuning strategies, balancing the benefits of massive pre-training against the need for domain specificity [18]. The system must manage a portfolio of models, each suited to different analytical windows: a fast, distilled model for real-time bias scoring on streaming data, and a larger, more expressive model for retrospective deep analysis of archived conversations. This heterogeneity introduces trade-offs between inference latency, throughput, and granularity that must be optimized according to the target application, whether high-frequency market surveillance or quarterly sentiment reporting. Big data principles further guide the storage layer, which must accommodate petabytes of unstructured text while enabling rapid retrieval for model fine-tuning cycles [19]. The unreasonable effectiveness of data, as observed in machine learning, implies that expanding the corpus with carefully curated investor narratives will likely improve bias detection accuracy, reinforcing the need for scalable data pipelines [20]. In addition, the architecture must incorporate feedback loops for model drift detection: when market regime shifts occur, the linguistic markers of biases may change, necessitating periodic recalibration.

4. Capturing Investor Self-Perception and Cognitive Biases

The extraction of self-perception from investor text moves beyond topic modeling and sentiment analysis to psychological profiling. Investor self-perception is conveyed through first-person singular and plural pronouns, modal verbs expressing certainty, comparative constructions, and retrospective evaluations of past decisions. For instance, a post claiming, “I knew this stock would bounce; my analysis was spot on—anyone who doubted me lacks patience,” reveals overconfidence and self-attribution bias simultaneously. An LLM fine-tuned on financial self-narratives can encode such linguistic patterns into latent representations that distinguish between calibrated confidence and excessive optimism. Fine-tuning strategies such as those investigated for text classification can be adapted to multi-label bias detection, where each sentence is tagged with bias categories including overconfidence, confirmation bias, loss aversion, and herding tendency [21]. The training data must be constructed through expert annotation, leveraging theoretical definitions from behavioral economics, and validated against external measures such as portfolio turnover and survey-based overconfidence scores.

The self-perception extraction module must engage with the insight that individuals construct their self-identity through discourse in ways that are both stable and context-sensitive [6]. Investors may present different facets of their self-concept depending on the online forum, the performance of their portfolio, and the social dynamics of the discussion thread. An LLM-based analytics system can capture this contextual variability by processing threads as coherent sequences rather than as isolated messages, using the surrounding conversation to disambiguate ironic boasts from genuine overconfidence. Moreover, contrastive learning objectives can be employed to teach the model to discriminate between narratives that reflect a dispositional bias versus situational attributions, enabling a more diagnostic measurement. By aggregating individual-level self-perception scores across thousands of investors, the system constructs time-varying indices of aggregate overconfidence or loss aversion, analogous to sentiment indices but rooted in cognitive mechanisms rather than mere optimism.

Beyond textual analysis, the system can incorporate multimodal signals—voice tone, facial expressions in video earnings calls—though these expand infrastructure demands substantially. Even within text alone, the granularity of LLM-driven bias measurement far exceeds what survey methods can achieve, because it taps into spontaneous, unprompted self-expression rather than retrospective self-reporting that is susceptible to social desirability biases. The resulting bias indicators can then be correlated with account-level trading records or anonymous portfolio data to validate their predictive power, closing the loop between linguistic self-representation and financial behavior.

5. Linking Self-Perception and Bias to Market Outcomes

The move from individual bias measurement to aggregate market outcomes requires explicit modeling of how distributed cognitive distortions coalesce into price dynamics. Classical behavioral finance theory postulates that overconfident investors trade more aggressively, contributing to excess volume and, when combined with short-sale constraints, to overpricing and subsequent crashes [7]. In the LLM analytics framework, aggregate overconfidence indices derived from social media narratives can be shown to lead trading volume and volatility at daily and intraday frequencies. Similarly, herding bias indicators, constructed from the prevalence of mimetic language and conformity expressions in forums, can help predict the timing of attention-driven rallies and subsequent reversals. The causal identification problem is formidable because market outcomes themselves feed back into investor narratives—rising prices may boost self-enhancing language, creating a reflexive loop. The system must therefore incorporate econometric identification strategies, such as instrumental variables based on exogenous platform outages or natural language processing shocks, to disentangle the bidirectional influences.

Case illustrations underscore the practical relevance. During the GameStop episode of early 2021, Reddit forums were replete with self-narratives that explicitly articulated a collective identity of retail investors taking on institutional short sellers. Language of certainty, defiance, and overestimated insight pervaded the discussions. An LLM-powered bias analytics pipeline would have detected a sharp rise in overconfidence and group-identity bias signals in the weeks prior to the price peak, alongside elevated herding language. Linked to real-time options and equity market data, such signals could provide early warnings of bubble formation. In cross-domain comparisons, similar patterns have been observed in cryptocurrency forums, where self-perception markers reliably precede boom-bust cycles. The architecture's capability to ingest heterogeneous text sources allows for the construction of cross-asset bias contagion maps, revealing how overconfidence propagates from one asset class to another through narrative spillovers. Big data methods for finance increasingly support such high-dimensional, cross-sectional analyses, enabling researchers to move from case studies to systematic surveillance [24]. Behavioral inattention models, which formalize how investors disregard certain signals, further enrich the interpretation of silence and absence of self-reflection in market narratives as an indicator of latent risk [25].

The integration of bias analytics with market microstructure data also permits the investigation of how self-perception-driven trading interacts with liquidity provision and volatility amplification at high frequencies. When a surge in overconfident language coincides with narrowing bid-ask spreads, the system may flag an impending liquidity mirage—a situation in which apparent depth evaporates as sentiment shifts. Thus, the LLM-powered analytics platform is not merely a sentiment dashboard but a comprehensive cognitive market monitor.

6. Governance, Fairness, and Robustness

Deploying LLM-based behavioral finance analytics at scale raises substantial governance challenges that span data ethics, fairness, transparency, and resilience against manipulation. The foremost concern is that the training corpora and model architectures may embed representational harms. If the textual data overrepresents certain demographic segments—for example, younger, male, English-speaking participants on Reddit—the resulting bias indices may systematically mischaracterize the cognitive profiles of underrepresented investor groups. This misalignment can lead to misguided policy interventions and reinforce existing inequities in financial market participation. Fairness frameworks developed for machine learning emphasize the need to audit both training data and model outputs for disparate performance across subpopulations; in the financial domain, such audits must consider intersectional dimensions including income, geography, gender, and language background [23]. The architecture must therefore include fairness monitoring modules that track subgroup bias index divergence and trigger recalibration when disparities exceed predefined thresholds.

Regulatory developments, most notably the European Union's proposed Artificial Intelligence Act, signal the growing demand for transparency, risk classification, and human oversight in AI systems used in critical sectors including finance [22]. An LLM analytics platform that influences investment decisions or market surveillance would likely fall under high-risk classification, requiring comprehensive documentation, conformity assessments, and post-market monitoring. The governance structure must incorporate model cards detailing training data provenance, intended use, known biases, and performance across stress scenarios. Additionally, the interpretability of LLM-based bias detection remains a frontier challenge. While attention weights and saliency maps can highlight which tokens drive a high overconfidence score, these explanations are often insufficient for regulatory purposes. The architecture should therefore combine neural models with symbolic reasoning components or retrieval-augmented generation that can produce human-readable justifications tied to canonical bias definitions.

Robustness to adversarial manipulation is another crucial dimension. Malicious actors could craft forum posts designed to suppress or exaggerate certain bias signals, misleading the analytics pipeline and potentially influencing market surveillance systems. Adversarial training, certification methods, and anomaly detection on incoming text distributions must be incorporated as defense layers. Moreover, the system must manage the risk of feedback loops wherein publicly disseminated bias indicators alter investor behavior, changing the very phenomena the analytics aim to measure. Continuous validation against out-of-sample market episodes and stress tests that simulate adversarial narrative campaigns are essential for maintaining trust in the analytics outputs.

7. Deployment and Sustainability Considerations

Operationalizing LLM-driven behavioral finance analytics requires careful orchestration of cloud infrastructure, model serving, and data lifecycle management. The streaming nature of financial conversations demands low-latency ingestion and inference, particularly when signals feed into algorithmic trading or real-time risk management. Microservice architectures, where separate containers handle text preprocessing, bias scoring, and market signal aggregation, enable horizontal scaling but introduce complexity in version synchronization and state management. The trade-off between using large, general-purpose foundation models and smaller, domain-specialized models is central. While foundation models offer superior few-shot capabilities and can capture subtle psychological nuances, their inference cost and carbon footprint can be prohibitive for high-frequency applications. Distillation, quantization, and architecture search can yield compact models that retain adequate detection accuracy for streaming use cases, aligning with sustainability goals [18]. The analytics platform should adopt a tiered strategy: a lightweight edge model processes high-volume streams in real time, flagging unusual activity that triggers deeper analysis by a larger model in a batch-oriented subsystem.

Model drift is a persistent concern. As market structures evolve and new financial instruments emerge, the linguistic conventions used by investors shift. The system must support continuous fine-tuning on recent data, with safeguards against catastrophic forgetting. Data pipelines should timestamp every ingested text and maintain versioned corpora for reproducible experiments. The energy consumption of repeated large-model training cycles demands efficiency measures such as curriculum learning, sparse expert models, and federated learning across institutions that hold sensitive investor data while preserving privacy. From a business continuity perspective, the deployment architecture must be resilient to cloud provider outages and cyberattacks, employing redundant instances and fallback models that can operate with degraded but still useful accuracy.

The sustainability of behavioral finance analytics also hinges on the long-term availability of high-quality textual data. As platforms change their APIs, restrict scraping, or shift toward more ephemeral content formats, the analytics pipeline must adapt by diversifying data sources and developing robust imputation methods. Finally, organizational sustainability requires building interdisciplinary teams that combine behavioral economists, NLP engineers, ethicists, and legal experts, ensuring that the system evolves in response to both technological advances and societal expectations.

8. Policy Implications

The deployment of LLM-based behavioral finance analytics has far-reaching implications for financial regulation and market oversight. Securities regulators could integrate such systems into market surveillance to detect nascent bubbles, coordinated manipulation campaigns, or destabilizing shifts in investor psychology before they manifest in price crashes. The fine-grained bias indicators offer a complement to traditional volatility and liquidity metrics, enabling a more preemptive macroprudential stance. However, the use of these tools by public authorities raises questions about proportionality and privacy: the same linguistic analysis that reveals collective overconfidence also exposes individual psychological profiles. Policy frameworks must delineate the boundary between aggregate surveillance and individual investor profiling, with strong anonymization and aggregation safeguards.

Another policy dimension involves disclosure and consumer protection. If investment platforms or robo-advisors begin to incorporate LLM-driven bias assessments into their recommendations, users must be informed about how their self-expressed cognitive traits influence the advice they receive. The fairness concerns noted previously demand that such systems not penalize investors for linguistic patterns associated with certain demographics, and that they provide avenues for contesting machine-generated bias characterizations. The EU AI Act and similar legislative efforts worldwide will likely mandate human-in-the-loop oversight for high-risk financial AI applications, requiring clear operational protocols [22]. Standard-setting bodies may need to develop industry-wide benchmarks for bias detection accuracy, transparency reports, and adversarial robustness testing, analogous to stress tests in banking. Cross-border coordination is essential because financial narratives and market contagion transcend jurisdictions, and inconsistent regulatory approaches could lead to regulatory arbitrage or fragmented market surveillance. In the long term, the integration of cognitive analytics into financial infrastructure could reshape the very nature of market efficiency, as reflexive awareness of bias metrics might prompt investors to correct their own heuristics—or, perversely, to strategically perform certain linguistic patterns to game the system. Policy must remain attentive to these evolutionary dynamics, fostering responsible innovation while safeguarding market integrity and investor welfare.

9. Conclusion

This paper has articulated a system-level framework for behavioral finance analytics that leverages large language models to explore the intricate relationship between investor self-perception, decision biases, and market outcomes. By moving beyond coarse sentiment measures to fine-grained extraction of self-referential narratives, the proposed architecture enables the construction of cognitive bias indices that can be empirically linked to trading volumes, volatility, and bubble dynamics. The analysis has highlighted the structural trade-offs inherent in such systems, including the tension between model expressiveness and inference cost, the challenges of ensuring fairness across heterogeneous investor populations, and the imperative of robustness against adversarial narrative manipulation. Governance frameworks derived from emerging AI regulations and fairness scholarship provide a scaffolding for responsible deployment, while sustainability considerations push toward efficient model architectures and federated data practices. Policy implications underscore the transformative potential of these analytics for market surveillance and investor protection, alongside the need for careful delineation of privacy, transparency, and oversight mechanisms. As financial markets become increasingly reflexive and information-rich, the ability to computationally decode the psychological underpinnings of investor behavior will not only deepen our scientific understanding but also enable the design of more resilient and equitable financial infrastructures. The continued cross-pollination between behavioral economics, natural language processing, and systems engineering will be critical to realizing this vision.

References

1. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263-291.

2. Baker, M., & Wurgler, J. (2007). Investor sentiment in the stock market. Journal of Economic Perspectives, 21(2), 129-151.

3. Shleifer, A. (2000). Inefficient markets: An introduction to behavioral finance. Oxford University Press.

4. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

5. Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. Journal of Finance, 66(1), 35-65.

6. Solanki, D., Hsu, H. M., Zhao, O., Zhang, R., Bi, W., & Kannan, R. (2020, July). The way we think about ourselves. In International Conference on Human-Computer Interaction (pp. 276-285). Cham: Springer International Publishing.

7. Odean, T. (1998). Volume, volatility, price, and profit when all traders are above average. Journal of Finance, 53(6), 1887-1934.

8. Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). Investor psychology and security market under- and overreactions. Journal of Finance, 53(6), 1839-1885.

9. Antweiler, W., & Frank, M. Z. (2004). Is all that talk just noise? The information content of internet stock message boards. Journal of Finance, 59(3), 1259-1294.

10. Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. Journal of Computational Science, 2(1), 1-8.

11. Chen, H., De, P., Hu, Y. J., & Hwang, B. H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. Review of Financial Studies, 27(5), 1367-1403.

12. Li, F. (2010). The information content of forward-looking statements in corporate filings—A naïve Bayesian machine learning approach. Journal of Accounting Research, 48(5), 1049-1102.

13. Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. Journal of Finance, 62(3), 1139-1168.

14. Jegadeesh, N., & Wu, D. (2013). Word power: A new approach for content analysis. Journal of Financial Economics, 110(3), 712-729.

15. Huang, A. H., Zang, A. Y., & Zheng, R. (2014). Evidence on the information content of text in analyst reports. The Accounting Review, 89(6), 2151-2180.

16. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. OpenAI blog, 1(8), 9.

17. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877-1901.

18. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.072Goldstein, I., Spatt, C. S., & Ye, M. (2021). Big data in finance. Review of Financial Studies, 34(7), 3213-3225.

19. Halevy, A., Norvig, P., & Pereira, F. (2009). The unreasonable effectiveness of data. IEEE Intelligent Systems, 24(2), 8-12.

20. Howard, J., & Ruder, S. (2018). Universal language model fine-tuning for text classification. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (pp. 328-339). Association for Computational Linguistics.

21. European Commission. (2021). Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). COM(2021) 206 final.

22. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (pp. 214-226). ACM.

23. Loughran, T., & McDonald, B. (2016). Textual analysis in accounting and finance: A survey. Journal of Accounting Research, 54(4), 1187-1230.

24. Gabaix, X. (2019). Behavioral inattention. In B. D. Bernheim, S. DellaVigna, & D. Laibson (Eds.), Handbook of Behavioral Economics: Applications and Foundations 2 (Vol. 2, pp. 261-343). North-Holland.