
What AI Analysis Can Reveal About Securities Class Actions

by [Mark Howrey](#) and [Emma Dong](#); Analysis Group, Inc.

Law360 (April 17, 2026)



Mark Howrey



Emma Dong

A [Law360](#) Pulse survey from March indicates that large numbers of attorneys are using artificial intelligence tools for an increasing variety of tasks, such as legal research and document creation.¹ Given this increased use of AI, one might also ask whether such tools can be helpful in undertaking data-driven assessments of litigation — particularly in the area of securities class actions.

In this article, we show how AI-based review of complaint text can enhance securities litigation analysis in two ways.

First, it enables more systematic identification of comparable cases by extracting a broader set of case characteristics. Second, it improves the ability to predict settlement amounts.² Using a database of hundreds of securities settlements, we find that measures of tone and certainty in complaints improve the accuracy of settlement predictions, particularly for larger cases.

These improvements are made possible by developments in AI, which allow for the extraction and analysis of large volumes of free-form text, such as complaints, or unstructured data. In comparison, data analysis has historically relied on information collected in standardized formats, or structured data, and manual review of limited amounts of unstructured data.

We also note the limits of what AI-assisted review can accomplish. In particular, certain important drivers of settlement outcomes, such as the level and structure of defendants' insurance policies, are not publicly observable. Thus, even with AI-based inputs, a large amount of variation in settlement outcomes remains unexplained.

Overview of AI-Based Analysis of Complaint Text

We used AI-assisted review of complaint text to extract structured information — e.g., the nature of allegations, and the channel and nature of the corrective disclosures — in more than 700 federal securities matters filed and settled after 2010.³

For example, we identified whether a case involves a short-seller report as the alleged corrective disclosure, or whether a government investigation is referenced in the complaint. In addition, we systematically measured multiple aspects of complaint language, such as tone and certainty,⁴ which would be challenging and time-consuming to evaluate through human review.

We also implemented human verifications to ensure that AI-generated outputs are properly refined before being incorporated into the statistical analysis.⁵ While AI can efficiently classify and extract information from large volumes of text, it may sometimes err on the side of inclusiveness when assigning categories, meaning that cases may be assigned to categories even when the supporting evidence is limited or ambiguous.

Therefore, we conducted multiple levels of validation, including parallel keyword searches and prompts designed to test the logical consistency of results. Any identified discrepancies were resolved by experienced professionals to ensure that the final classifications accurately reflected the relevant case characteristics.⁶

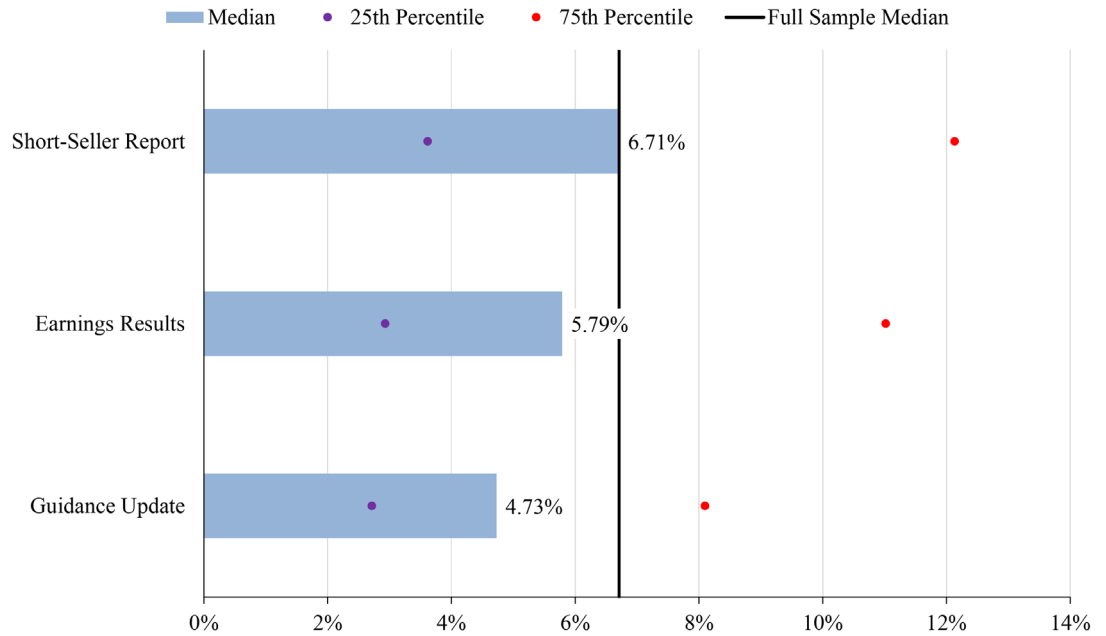
Beyond data verification, it is essential to examine whether the results of any AI-assisted analysis make economic sense. As described below, the results of our analysis of comparable cases, based on a combination of AI-generated data and human verification, have intuitive appeal.

Example Application: Extraction of Case Characteristics for Comparison

First, we analyzed how settlement outcomes differ by the channel and nature of the corrective disclosure, when the outcome is measured as a percentage of simplified plaintiff-style damages — the settlement percentage.⁷

The figures below compare the distribution of settlement percentages for different classifications to the median for the full sample of over 700 settlements.

Figure 1: Settlement Percentage by Channel of Last Corrective Disclosure



For example, in Figure 1, our analysis suggests that the channel of the last corrective disclosure,⁸ such as whether it was revealed in a short-seller report, has minimal impact on the settlement percentage, with a median settlement percentage for that classification being nearly the same as the median for the full sample. In contrast, when the corrective disclosure consists solely of an earnings miss or a guidance update,⁹ without additional substantive revelations, settlement percentages tend to be lower.

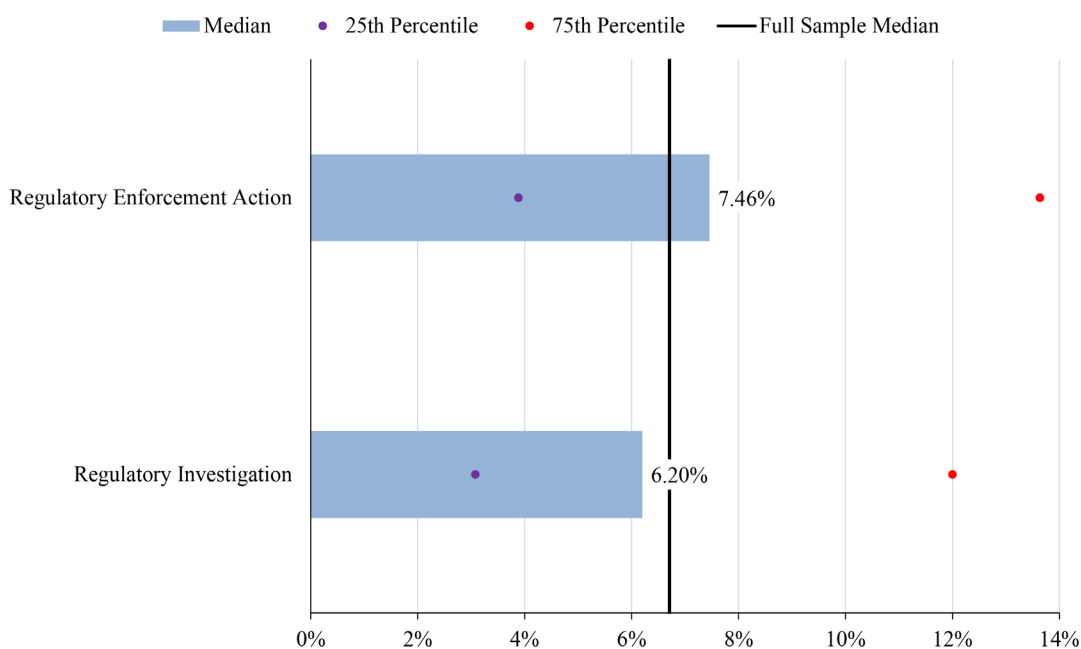
This is economically intuitive: Companies may miss earnings or adjust guidance, which is forward-looking in nature, for many reasons other than fraud.

Similarly, differences in the nature or subject matter of the information revealed in the corrective disclosures are associated with different settlement percentages, as shown in Figure 2.

For example, cases in which the corrective disclosure references the existence of a regulatory investigation, without further action, are associated with lower settlement percentages compared to cases in which the corrective disclosure reveals an enforcement action.¹⁰

This distinction is also economically intuitive: An investigation alone may reflect uncertainty, whereas an enforcement action may signal more concrete evidence of any alleged misconduct.

Figure 2: Settlement Percentage by Nature of Last Corrective Disclosure



As shown in Figures 1 and 2, substantial variation in settlement percentages remains within each category, which reflects the impact of unobserved case-specific factors, such as the strength of the evidence and insurance considerations. While we provide a few examples of information that can be extracted from an AI-assisted review of complaint text, many more characteristics could be identified and analyzed.

Example Application: Sentiment Analysis

We also examined whether the tone and certainty reflected in the complaints help explain settlement amounts. Sentiment analysis is a well-known method for systematically extracting signals in language that may not be explicitly stated.¹¹

Without the assistance of AI, measuring sentiment is time-consuming, often requiring multiple human reviews of the underlying documents and a process to reconcile differences in scoring. AI allows this classification to be performed efficiently, and similar approaches have been used in academic research.¹²

Following established sentiment frameworks,¹³ we used AI to categorize complaint language along two primary dimensions: the level of certainty (e.g., confident versus qualified language) and emotional tone, which includes negativity, intensity (e.g., calm versus excited language), and the extent to which the complaint uses assertive and definitive language.¹⁴

Conceptually, these measures may capture certain so-called soft information about the perceived strength of plaintiffs' claims. For example, when plaintiffs believe they

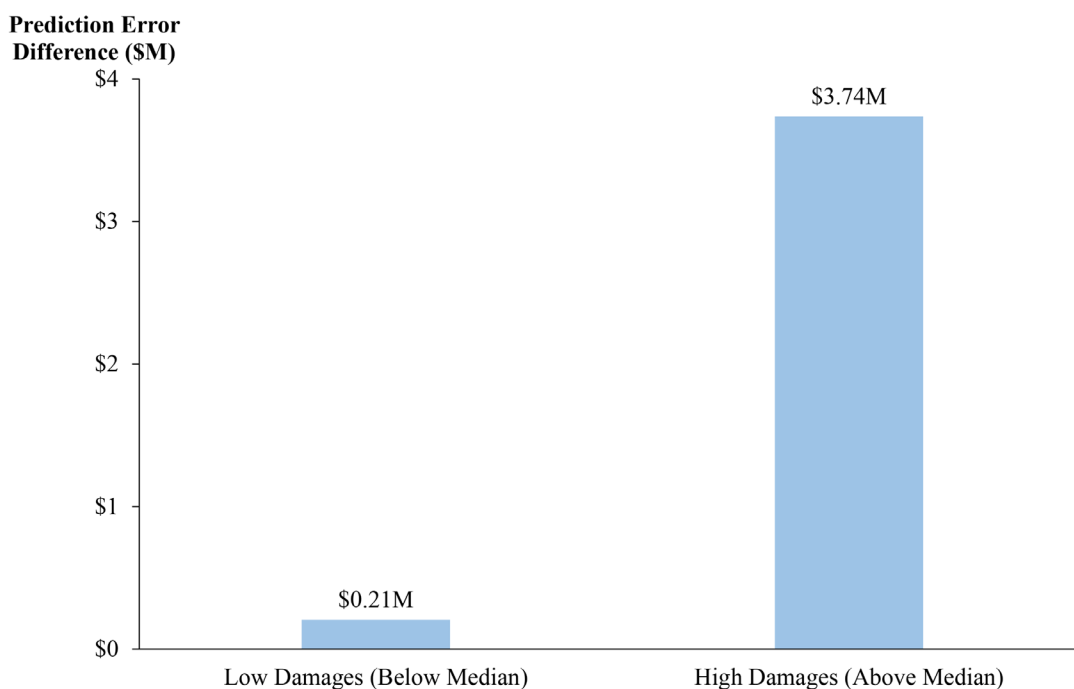
have stronger evidence, complaints may be more direct and confident, with fewer qualifiers or speculative phrases.

To examine whether the sentiment information from the complaints helps explain settlement amounts, we compared the prediction errors of a model that predicts settlements using only plaintiff-style damages — i.e., the benchmark model¹⁵ — against those of a model that includes the sentiment variables, i.e., the sentiment model.¹⁶

As shown in Figure 3, the sentiment model reduces prediction errors, and the impact is greater for larger settlements.¹⁷

To put this in context, the reduction in prediction error represents roughly 5% of the median settlement in smaller cases (\$210,000 out of \$3.98 million) and about 12% of the median settlement in larger cases (\$3.7 million out of \$31.64 million). For larger cases, the \$3.7 million average reduction means that prediction errors are approximately 8% smaller than under the benchmark model.

Figure 3: Reduction in Prediction Error From the Benchmark Model to the Sentiment Model



One possible explanation for the difference in explanatory power is that smaller cases often settle at amounts influenced by considerations such as fixed litigation costs, whereas larger cases may be negotiated more closely around the merits, so differences in language sentiment can better explain variations in settlement outcomes beyond the size of damages.

However, complaint text does not directly capture factors that may affect settlement outcomes, including the amount and structure of defendants' insurance. AI-assisted review should therefore be viewed as a complement to, not a substitute for, legal and economic judgment.

Conclusion

When AI-based reviews of complaint text are combined with economic expertise and professional judgment, they can enhance securities class action analysis in two key ways: by enabling more systematic identification of comparable cases and by improving the accuracy of settlement predictions, particularly in larger cases, while some variation in settlement amount remains unexplained due to unobserved factors.

Mark Howrey is a managing principal and [Emma Xiaoxiao Dong](#) is a vice president at Analysis Group, Inc.

Analysis Group managing principal [Gaurav Jetley](#) contributed to this article.

The opinions expressed are those of the author(s) and do not necessarily reflect the views of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

Endnotes

- 1 Kevin Penton, "What Attorneys Really Think About AI," Law360, March 31, 2026, available at <https://www.law360.com/corporate/articles/2454609>.
- 2 As discussed in the article "Assessing Federal Securities Class Action Stats In '25 So Far," AI-assisted review can also be used to identify the relationship between settlement amounts and the litigation stage of the case at the time of settlement. Howrey, Mark, and Emma Dong, "Assessing Federal Securities Class Action Stats in '25 So Far," Law360, August 21, 2025, available at: <https://www.law360.com/articles/2376958/assessing-federal-securities-class-action-stats-in-25-so-far>.
- 3 The sample includes federal class action settlements with Rule 10b-5 or Securities Act Section 11 claims, excluding settlements related to mergers and acquisitions, those brought by the [U.S. Securities and Exchange Commission](#), and those where the alleged misrepresentations or omissions were associated with securitizations. Approximately 5% of cases were excluded from the analysis because the complaint text could not be reliably processed.
- 4 Methods for extracting language sentiment, including the design of AI prompts, may vary depending on the type of text being analyzed. For example, prompts for shorter informal content such as social media posts may differ from those used for longer formal documents such as legal complaints.
- 5 The importance of human-in-the-loop systems is documented, especially for tasks requiring context-sensitive judgment and nuanced interpretation. See "Human in the Loop AI: Elevating Accuracy, Oversight, and Trust in AI Systems," WitnessAI, November 10, 2025, available at: <https://witness.ai/blog/human-in-the-loop-ai/>.
- 6 For example, the human-in-the-loop process helps ensure that classifications align with how practitioners and economists distinguish between types of disclosures (e.g., government investigations versus government enforcement actions) is consistent with the AI results.

- 7 Plaintiff-style damages are calculated using the inflation-based per-share damages specified in the settlement plan of allocation. Trading behavior is estimated using a two-trader model. For benchmarking purposes, shares available to trade and trading volume are not adjusted for short interest or insider holdings. While estimating trading behavior using a trading model is frequently used in settlement negotiations, courts have rejected the approach as speculative.
- 8 We use an AI-based method that reviews the complaints and identifies natural groupings based on similarities in the channel and nature of the corrective disclosures. Rather than forcing cases into predetermined classifications, this approach allows the possible categories to emerge from the data itself. The generated list is then manually reviewed and adjusted.
- 9 An “earnings result” disclosure refers to a company-issued earnings release or preannouncement that allegedly reveals adverse financial or operational information. A “guidance update” refers to an announcement revising forward-looking guidance that was allegedly based on a revelation of previously concealed facts. Cases with a corrective disclosure that involved an earnings miss and a revised guidance are included in both categories.
- 10 A classification of “Regulatory Investigation” involves disclosures of government investigations (such as subpoenas), whereas a classification of “Regulatory Enforcement Action” requires disclosures of an enforcement action, such as consent orders, fines, or settlements. For example, a company disclosure of an SEC subpoena regarding its accounting practices is considered a “Regulatory Investigation.” A company disclosure that it entered a consent order with the FDIC is considered a “Regulatory Enforcement Action.”
- 11 See Tan, Kian L., et al. “A Survey of Sentiment Analysis: Approaches, Datasets, and Future Research,” *Applied Sciences*, 2023, Vol. 13(7), p. 4550 and Kumar, Mahander, et al., “Evolving Techniques in Sentiment Analysis: A Comprehensive Review,” *PeerJ Computer Science*, 2025, e2592 for reviews of the applications and methodologies of sentiment analysis.
- 12 For example, see Jiang, Wei, et al. “Measuring Sustainability with AI,” *Artificial Intelligence, Finance, and Sustainability*, 2024, Chapter 3, pp. 33–57, for a review of how AI is used to measure firms’ ESG exposure, including employing sentiment analysis. In legal contexts, AI-powered sentiment analysis has been used to unearth hidden biases within case law and to examine the semantic structures of bills. See Abimbola, Bolanle, et al. “Enhancing Legal Sentiment Analysis: A Convolutional Neural Network–Long Short-Term Memory Document-Level Model,” *Machine Learning and Knowledge Extraction*, 2024, Vol. 6(2), pp. 877–897; Valdez, Danny, and Patricia Goodson. “Neutral or Framed? A Sentiment Analysis of 2019 Abortion Laws,” *Sexuality Research and Social Policy*, 2022, Vol. 19, pp. 936–945
- 13 Hyland, Ken. “Boosting, Hedging and the Negotiation of Academic Knowledge,” *Text & Talk*, 1998, Vol. 18(3), pp. 349–382. See also Daniel, Florence O. “Dissenting with Conviction and Deference: Boosters and Hedges in Nigerian Supreme Court Dissents,” *International Journal for the Semiotics of Law*, Vol. 38, 1579–1602. Broekens, Joost, et al., “Fine-grained Affective Processing Capabilities Emerging from Large Language Models,” 11th International Conference on Affective Computing and Intelligence Interaction, 2023.
- 14 AI classifies sentiment at the sentence level. We then assign numeric values for each sentiment category and take an average across all sentences in the complaint.
- 15 Plaintiff-style damages alone are known to be a driver of settlement amounts, even without making additional modeling assumptions, so the benchmark regression model includes $\log(\text{damages})$ and its squared term.
- 16 In the sentiment model, we include indicator variables for common plaintiff law firms to control for potential firm-specific language style effects that may influence measured sentiment.
- 17 To evaluate prediction error, we conduct 10,000 simulations in which 80% of the cases are randomly selected to train the model and the remaining 20% are used for testing. In each iteration, we calculate the average absolute dollar difference between predicted and actual settlements for the test cases (“Dollar MAE”). The reported prediction error is the median Dollar MAE across all iterations. In this article, we define “larger cases” as those with damages above the sample median, although the results are similar if case size is defined using settlement amounts.