FINANCIAL NLP CASE STUDY — APPROACH & RESULTS

Tesla vs. Rivian

Comparing Language in 10-K Filings & Earnings Calls with NLP

Sentiment Analysis | Readability Scoring | Year-over-Year Similarity | Topic Extraction

Data: 10-K Item 1A & Item 7 (FY2023–FY2024) · Earnings call transcripts · CEO/CFO prepared remarks

Objective

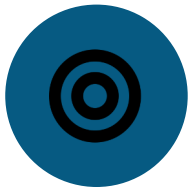

What does each company's own language reveal about how it frames risk, performance, and outlook — and how does that framing shift year over year?

Two Companies

Tesla (TSLA) and Rivian (RIVN) — an established EV leader vs. a newer challenger, compared on equal footing.

Two Document Types

Regulatory 10-K filings (formal, legally reviewed) vs. earnings call transcripts (spontaneous, investor-facing).

Two Fiscal Years

FY2023 vs. FY2024 10-K sections, to see how each company's own language evolved year over year.

Data Sources

Three text sources per company, spanning formal filings to spontaneous commentary

10-K Filings

Item 1A (Risk Factors) and Item 7 (MD&A), for both FY2023 and FY2024 — extracted directly from each company's PDF filing by page range.

Earnings Call Transcripts

Most recent quarterly call for each company, split into the Management Discussion section and the Q&A section.

CEO / CFO Prepared Remarks

Standalone written comments from each company's CEO and CFO, analyzed separately to compare leadership tone.

Analysis Pipeline

Every document moves through the same four stages before comparison

1 Extract

Pull raw text from PDFs (10-Ks, transcripts) by page range or section marker.

2 Clean

Multi-pass regex cleaning: HTML/markdown artifacts, boilerplate, speaker labels, whitespace.

3 Analyze

Score each document for sentiment, readability, similarity, and topics.

4 Compare

Line results up across company and fiscal year to surface differences.

Method — Sentiment Analysis

Two independent, complementary scoring approaches applied to every text source

Loughran-McDonald (LM) Dictionary

Finance-specific lexicon

- Built specifically for 10-Ks and other financial disclosures — avoids the false positives general dictionaries produce on financial language
- Scores each document across seven categories: positive, negative, uncertainty, litigious, strong modal, weak modal, and constraining
- Category scores are reported as a percentage of total dictionary-matched words

Bing Liu Opinion Lexicon

General-purpose sentiment + negation

- A general positive/negative word list (via NLTK), used as an independent check against the finance-specific LM score
- Produces a single averaged sentiment score per document
- Paired with a negation-word count, since 'not strong' and 'strong' should not score identically

Method — Readability Scoring

Gunning Fog Index, applied to each company's earnings call transcript

$$\text{Fog} = 0.4 \times \left(\frac{\text{words}}{\text{sentences}} + 100 \times \frac{\text{complex words}}{\text{words}} \right)$$

Higher score = harder to read. A word counts as "complex" if it has 3 or more syllables.

- **Transcripts get an extra cleaning pass first**

removes speaker/title labels (e.g. "Elon Musk, Chief Executive Officer"), timestamps, and boilerplate that would otherwise distort sentence and word counts

- **Syllables are estimated with a vowel-cluster heuristic**

a lightweight approximation rather than a phonetic dictionary lookup — fast, but not exact

Why it matters:

Readability is used as a proxy for communication style — a persistently high Fog score can suggest denser, more technical language, while a lower score suggests plainer, more direct communication with investors.

Method — Year-over-Year Similarity

How much did the wording of each 10-K section change between FY2023 and FY2024?

1. Vectorize

TF-IDF vectors built per 10-K section, across both companies and both years, with common stopwords removed

2. Score

Cosine similarity computed between each company's 2023 and 2024 vector, for Item 1A and Item 7 separately

3. Interpret

Score near 1.0 = little year-over-year change in wording; score near 0 = substantial rewording

Four comparisons produced per fiscal-year pair:

- Tesla — Item 7
- Tesla — Item 1A
- Rivian — Item 7
- Rivian — Item 1A

Method — Topic Extraction

TF-IDF keyword ranking, applied to 10-K sections and earnings calls alike

- Each document is ranked by TF-IDF weight; the top-scoring words serve as a lightweight proxy for its key topics
- A custom stopword list — combining standard English stopwords, a common-word frequency list, and manually identified finance/legal boilerplate — was built to filter out generic terms (“company,” “revenue”) and let more distinctive language surface
- Topics from 2023 and 2024 are then compared per company and section

Added

New topics that appear only in 2024

Removed

Topics present in 2023 but dropped by 2024

Unchanged

Topics that persisted across both years

Results — Executive Summary

Four headline findings from running the full pipeline on Tesla and Rivian

0.0 vs +1.3 pts

Net tone, full earnings call

Tesla's call was sentiment-neutral overall; Rivian's skewed net positive.

12.2 vs 14.9

Gunning Fog readability

Rivian's call used denser, harder-to-read language than Tesla's.

9/10 vs 3/10

Item 7 topics unchanged YoY

Tesla's MD&A stayed on similar topics; Rivian's shifted the most.

FSD & Autonomy vs EBITDA & Capex

Top earnings-call keywords

Tesla's call leaned into a tech/autonomy narrative; Rivian's leaned into financial fundamentals.

Results — Sentiment by Section

Net tone (LM percent-positive minus percent-negative) across three call sections

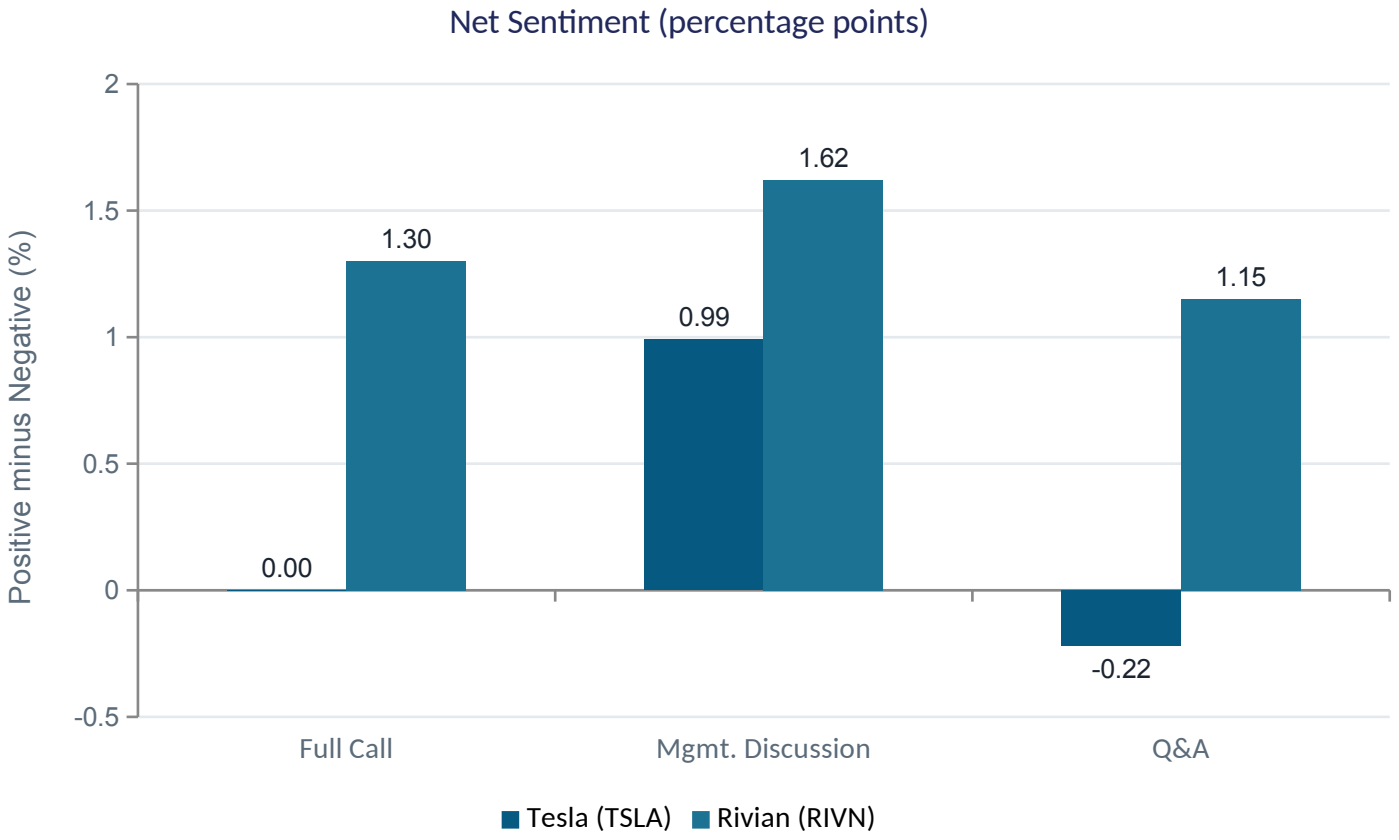

What stands out

- Tesla's Q&A is the only net-negative section in the whole analysis
- Tesla's Q&A also has 89 negation-word hits vs. 31 for Rivian — far more “not / don't / can't”-style hedging under analyst questioning
- Rivian's net tone stays positive in every section, including Q&A

Results — Readability & Leadership Tone

Gunning Fog Index — full earnings call

Higher = harder to read. Both scores sit in the "fairly difficult" range typical of executive speech.

Net Sentiment — CEO vs. CFO Remarks

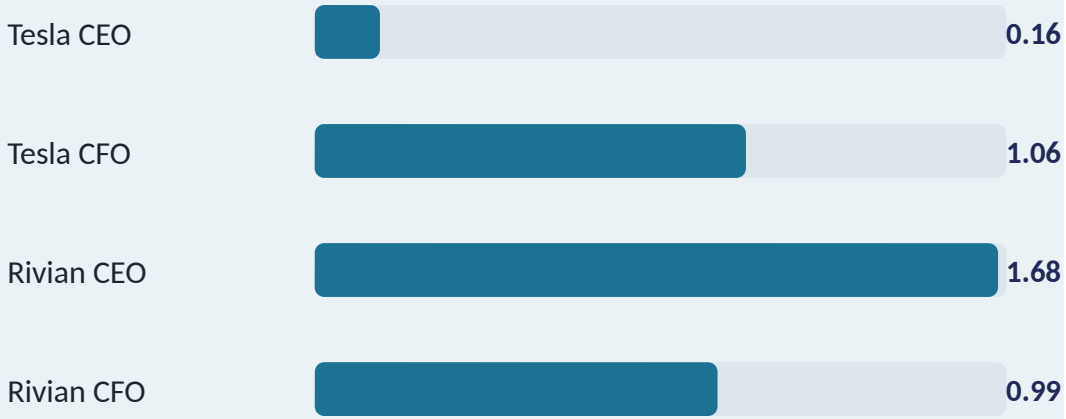

Rivian's CEO used the most net-positive language of any single speaker; Tesla's CEO the least.

Results — Year-over-Year Stability

2023 vs. 2024 10-K sections: cosine similarity and how many of the top 10 topics persisted

Cosine Similarity (TF-IDF)

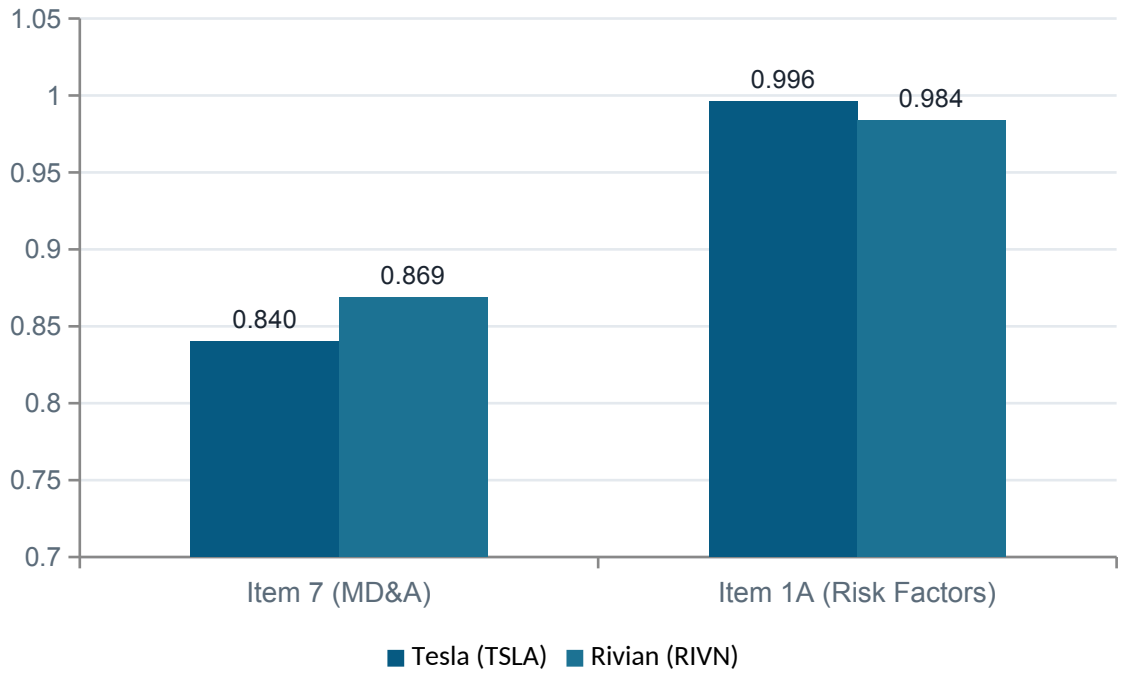

Top-10 Topics Unchanged Year-over-Year

Risk Factors (Item 1A) barely change either year — for both companies. Item 7 (MD&A) moves more, and Rivian's shifted the most: only 3 of its top 10 words carried over.

Results — What Each Company Talked About

Top TF-IDF keywords from each company's most recent earnings call

Tesla — a tech & autonomy narrative

technoking
optimus
fsd
unsupervised
autonomy
lidar
autonomous
axelrod
vaibhav
taneja

Rivian — a financial-fundamentals narrative

autonomy
cogs
capex
ebitda
gaap
incremental
learnings
gpus
genuity
gianarikas

Reading the contrast: Tesla's call centers on Full Self-Driving, Optimus, and autonomy — a product/technology story. Rivian's centers on COGS, CapEx, EBITDA, and GAAP — a cost-discipline and path-to-profitability story. "Autonomy" is the one word both calls share.

Tools & Technical Stack

Text Extraction

- PyPDF2 — PDF page reading & text extraction
- html / re (regex) — markup and pattern cleaning

Sentiment & Language

- Loughran-McDonald Master Dictionary
- NLTK — opinion lexicon, tokenization, negation list

Vectorization & Similarity

- scikit-learn TfidfVectorizer
- scikit-learn cosine_similarity

Stopwords & Vocabulary

- spaCy English stopwords
- NLTK stopwords + curated finance/name terms

Key Takeaways

Across every method, the same pattern holds: Rivian communicates more consistently positive, Tesla shows more variation by section.

1. Tesla's tone is section-dependent — net-neutral on the full call, net-positive in prepared remarks, but the only net-negative result anywhere in this analysis shows up in Tesla's Q&A, alongside nearly 3x the negation-word count of Rivian's Q&A

2. Rivian's tone is consistently net-positive across every section and speaker, and its CEO's remarks carried the most net-positive language of any single speaker measured

3. 10-K Risk Factors (Item 1A) barely move year-over-year for either company (cosine similarity ≥ 0.98); Item 7 (MD&A) moves more — especially for Rivian, where only 3 of 10 tracked topics carried over from 2023 to 2024

4. The two earnings calls tell different stories in their own words: Tesla's centers on autonomy/FSD/Optimus, Rivian's centers on COGS/CapEx/EBITDA — a tech narrative vs. a financial-discipline narrative

Caveats: 10-K page ranges were located manually per filing; syllable counts for the Fog index are a heuristic, not a dictionary lookup; sentiment/readability are proxies, not direct measures of company performance.