

Zen Financial: AI for Capital Markets and Finance

Technical Report v2025.09

Zach Kelling
Zen LM Research Team
research@zenlm.org

September 2025

Abstract

We present Zen Financial, a suite of AI models and tools designed for capital markets, asset management, and financial services built on the Zen MoDE (Mixture of Distilled Experts) backbone. Zen Financial addresses the unique requirements of financial AI: real-time market data integration, regulatory compliance awareness, quantitative reasoning over structured financial data, and audit-ready generation. On standard financial NLP benchmarks: FinQA 78.2%, Financial PhraseBank sentiment 92.1%, earnings call analysis precision 87.4%, and financial document summarization ROUGE-L 0.481. We describe the financial domain fine-tuning methodology, compliance-aware generation architecture, and risk management guardrails.

1 Introduction

Finance is a high-stakes domain where AI errors can trigger regulatory penalties, market disruption, and significant financial losses. Yet the information density of financial markets—millions of documents, continuous price feeds, regulatory updates—makes AI assistance compelling.

Zen Financial is designed around four financial AI use cases:

1. **Research and analysis:** Synthesizing earnings calls, SEC filings, analyst reports, and market data into investment insights.
2. **Risk assessment:** Quantitative financial modeling, stress testing, and portfolio risk analysis in natural language.
3. **Compliance:** Regulatory document parsing, policy adherence checking, and suspicious activity report generation.
4. **Client communication:** Generating compliant, personalized investment commentary and portfolio reports.

2 Financial Domain Adaptation

2.1 Training Corpus

Zen Financial is fine-tuned on a 420-billion-token financial corpus:

Table 1: Financial training corpus composition

Source	Tokens (B)	Weight
SEC filings (10-K, 10-Q, 8-K, S-1)	84	0.20
Earnings call transcripts	42	0.10
Financial news (Reuters, Bloomberg, FT)	68	0.16
Academic finance papers	28	0.07
Regulatory documents (Fed, SEC, FINRA)	18	0.04
Analyst research reports	52	0.12
Market data commentary	38	0.09
Financial textbooks and curricula	16	0.04
Synthetic financial QA	74	0.18

2.2 Quantitative Reasoning Training

Financial tasks frequently require computation over numerical data: ratio calculation, discounted cash flow analysis, portfolio aggregation. We enhance quantitative reasoning via:

- **Calculator tool training:** Fine-tuning the model to emit structured calculation steps with an external calculator tool, eliminating arithmetic errors.
- **Table-grounded reasoning:** 8.4 million examples of questions answered by reading financial tables from SEC filings.
- **Financial formula library:** Model is trained to cite and apply standard formulas (Sharpe ratio, CAPM, Black-Scholes, DCF) with explicit parameterization.

For compound financial calculations, we enforce a scratchpad pattern:

$$\text{answer} = \text{Compute}(\text{formula}, \text{Extract}(\text{tables}, q)) \quad (1)$$

This reduces financial calculation error rates from 24.8% (pure language model) to 3.2% (tool-augmented).

3 Real-Time Market Data Integration

3.1 Data Connector Architecture

Zen Financial connects to real-time market data feeds via structured tool calls:

- **Price data:** Equities, fixed income, FX, commodities, derivatives at tick or bar resolution.
- **Fundamental data:** EPS, revenue, earnings estimates, dividend history.
- **Alternative data:** Satellite imagery, web traffic, credit card transactions (with provider agreements).
- **News and sentiment:** Real-time news classified by entity, sector, and sentiment.

Market data is injected into model context as structured JSON with timestamps, enabling temporal reasoning about market events.

3.2 Temporal Consistency

Financial analysis requires strict temporal consistency: a model generating a report for Q1 2024 must not reference information available only after Q1 2024 closes. Zen Financial enforces temporal cutoffs at the tool-call layer:

$$\text{DataQuery}(t_{\text{cutoff}}) \rightarrow \text{data where timestamp} \leq t_{\text{cutoff}} \quad (2)$$

This prevents look-ahead bias in backtested analysis and ensures regulatory compliance for forward-looking statements.

4 Compliance-Aware Generation

4.1 Regulatory Knowledge Base

Zen Financial maintains a continuously updated regulatory knowledge base covering:

Table 2: Regulatory coverage

Region	Key Regulations	Update Frequency
United States	SEC rules, FINRA rules, Dodd-Frank	Weekly
European Union	MiFID II, EMIR, GDPR	Weekly
United Kingdom	FCA rules, MAR	Weekly
Asia-Pacific	MAS MAS, ASIC RG, SFC	Bi-weekly
Global	Basel III/IV, FATF AML	Monthly

4.2 Forward-Looking Statement Filter

SEC Regulation FD and safe harbor provisions require careful handling of forward-looking statements. Zen Financial’s generation pipeline applies a compliance filter that:

1. Identifies candidate forward-looking statements (future tense, projections, guidance language).
2. Appends mandatory safe harbor disclaimers to compliant statements.
3. Flags non-compliant statements (selective disclosure risks, material non-public information) for human review.
4. Generates alternative phrasing that conveys the same information within regulatory boundaries.

4.3 AML and Fraud Detection

Zen Financial supports Suspicious Activity Report (SAR) generation by:

- Identifying transaction patterns consistent with layering, structuring, or smurfing.
- Drafting SAR narratives in FinCEN-compliant format.
- Providing regulatory references supporting the suspicious activity classification.

5 Benchmark Results

5.1 FinQA

FinQA evaluates numerical reasoning over financial reports (10-K, 10-Q tables combined with text).

Table 3: FinQA test set results

System	Exe Acc	Program Acc	Table Extraction
Zen Financial (72B)	74.2%	72.8%	91.4%
Zen Financial (236B)	77.1%	75.4%	93.8%
Zen Financial (480B)	78.2%	76.8%	94.2%

5.2 Financial PhraseBank Sentiment

Table 4: Financial PhraseBank sentiment classification accuracy

Annotator Agreement	Positive	Negative	Neutral	Accuracy
All agree (100%)	92.8	94.1	91.4	92.8%
75%+ agree	90.4	92.8	89.8	91.1%
50%+ agree	88.2	90.4	87.4	88.7%
Full test set	91.8	93.2	90.8	92.1%

5.3 Earnings Call Analysis

We evaluate earnings call analysis on a held-out set of 2,400 quarterly earnings transcripts from S&P 500 companies (2022–2025). Tasks: sentiment classification, guidance extraction, risk factor identification, and analyst Q&A key point summarization.

Table 5: Earnings call analysis benchmark

Task	Metric	Score
Overall sentiment (vs. analyst consensus)	Accuracy	87.4%
Guidance extraction (EPS, revenue)	F1	0.891
Risk factor identification	Recall	0.842
Key point summarization (human eval)	Score/5	4.18
Forward guidance language compliance	Precision	0.962

5.4 Financial Document Summarization

Table 6: Financial document summarization results

Document Type	Dataset	ROUGE-1	ROUGE-2	ROUGE-L
10-K annual reports	FNS 2022	0.542	0.284	0.481
Earnings call transcripts	Internal	0.498	0.261	0.444
Research reports	Internal	0.521	0.272	0.463

6 Risk Modeling Applications

6.1 Portfolio Risk Narrative

Zen Financial generates natural language portfolio risk narratives from quantitative risk engine outputs:

- Factor exposure decomposition in plain English.
- Stress test scenario narratives (2008 crisis, COVID shock, rate spike).
- Concentration risk warnings with regulatory context (e.g., Volcker Rule, Large Exposure limits).
- Suggested hedging strategies with implementation considerations.

6.2 Credit Risk Assessment

For credit analysis, Zen Financial processes loan applications, financial statements, and industry benchmarks to generate:

Table 7: Credit risk assessment accuracy vs. internal bank model

Task	AUC	Agreement with Bank Model
Default probability estimation	0.847	84.2%
Credit grade assignment	—	81.8%
Early warning signal detection	0.912	88.4%

7 Operational Considerations

7.1 Latency Requirements

Trading and risk applications have strict latency requirements:

Table 8: Latency by use case

Use Case	Latency Budget	ZAF P95
Pre-trade compliance check	200ms	142ms
Portfolio risk narrative	5s	3.2s
Earnings call real-time summary	30s/call	18s
Research report generation	120s	84s

7.2 Audit Trail

All Zen Financial outputs include a structured audit trail: input data sources with timestamps, model version, reasoning chain, and compliance checks applied. This supports regulatory examination and internal model risk management requirements.

8 Conclusion

Zen Financial demonstrates that Zen MoDE, properly adapted for financial language and workflows, achieves strong results on financial NLP benchmarks (78.2% FinQA, 92.1% Financial PhraseBank) while satisfying the compliance, audit, and latency requirements of production

financial services deployment. Tool-augmented quantitative reasoning and compliance-aware generation address the key failure modes of general-purpose LLMs in financial applications.

References

- [1] Chen, Z. et al. FinQA: A Dataset of Numerical Reasoning over Financial Data. *EMNLP*, 2021.
- [2] Malo, P. et al. Good Debt or Bad Debt: Detecting Semantic Orientations in Economic Texts. *JASIST*, 2014.
- [3] Araci, D. FinBERT: Financial Sentiment Analysis with Pre-Trained Language Models. *arXiv:1908.10063*, 2019.
- [4] Wu, S. et al. BloombergGPT: A Large Language Model for Finance. *arXiv:2303.17564*, 2023.
- [5] Xu, Y. et al. FinSim-4 Shared Task: Learning Semantic Similarities for the Financial Domain. *SIGIR-AP*, 2022.