

# 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 domain-adapted language model and tool suite for capital markets, asset management, and financial services. Zen Financial is *not* a from-scratch model: it is a supervised fine-tune of the openly licensed **Qwen3-8B** base model (developed by Alibaba and released under the Apache-2.0 license), adapted to the financial domain using the Zen-Pro fine-tuning recipe. Building on Qwen3-8B’s general reasoning and instruction-following capabilities, Zen Financial targets the unique requirements of financial AI: real-time market data integration via tool calls, regulatory compliance awareness, quantitative reasoning over structured financial data, and audit-ready generation. We describe the financial domain fine-tuning methodology, the compliance-aware generation architecture, and the risk-management guardrails. We deliberately do not report headline accuracy figures on financial NLP benchmarks here; instead we describe the adaptation qualitatively and name the public benchmarks (FinQA, Financial PhraseBank, and others) against which such models should be evaluated under transparent, reproducible protocols.

## 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 Base Model and Licensing

Zen Financial is a domain fine-tune, not a model trained from scratch. The base model is **Qwen3-8B** [1], an 8-billion-parameter dense transformer developed and released by Alibaba

under the Apache-2.0 license. We selected Qwen3-8B for its strong general reasoning, tool-use behavior, and long-context handling, together with a permissive license that allows redistribution of derived weights. Every capability described in this report is the result of *adapting* this base model to the financial domain; the underlying architecture, pre-training, and general-domain knowledge are inherited from Qwen3-8B and credited to its original authors. Zen Financial is *not* based on a bespoke “Zen MoDE” or mixture-of-experts backbone, and we describe no multi-hundred-billion-parameter model family.

## 2.2 Fine-Tuning Corpus

Zen Financial is fine-tuned on a financial corpus assembled from the sources below. This corpus drives continued pre-training and supervised fine-tuning on top of Qwen3-8B; the *Share* column gives sampling proportions during adaptation, not pre-training-scale token budgets.

Table 1: Financial fine-tuning corpus composition (sampling share during domain adaptation of Qwen3-8B)

| Source                                  | Share |
|-----------------------------------------|-------|
| SEC filings (10-K, 10-Q, 8-K, S-1)      | 0.20  |
| Earnings call transcripts               | 0.10  |
| Financial news (Reuters, Bloomberg, FT) | 0.16  |
| Academic finance papers                 | 0.07  |
| Regulatory documents (Fed, SEC, FINRA)  | 0.04  |
| Analyst research reports                | 0.12  |
| Market data commentary                  | 0.09  |
| Financial textbooks and curricula       | 0.04  |
| Synthetic financial QA                  | 0.18  |

## 2.3 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:

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

Delegating arithmetic to an external calculator tool, rather than relying on the language model to compute in-context, substantially reduces financial calculation errors; the scratchpad pattern isolates the formula and inputs so that the numeric step is deterministic and auditable. We report no specific error-rate figure here.

## 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 Evaluation Methodology

Rather than report headline scores, this section names the public benchmarks against which a financial language model such as Zen Financial should be evaluated, and the protocol we recommend. Because Zen Financial is a fine-tune of Qwen3-8B, the meaningful comparison is between the adapted model and the unmodified Qwen3-8B base, with the evaluation harness and prompts held fixed, so that any measured difference is attributable to the financial adaptation rather than to a different or larger backbone.

### 5.1 Recommended Public Benchmarks

- **FinQA** [2]: numerical reasoning over financial reports, combining 10-K/10-Q tables with text. This is the primary public test of the tool-augmented quantitative reasoning that the adaptation targets, and should be reported with the dataset’s standard execution-accuracy and program-accuracy metrics.
- **Financial PhraseBank** [3]: sentence-level financial sentiment classification, with results conventionally stratified by annotator-agreement level.
- **FinBERT** [4] and **BloombergGPT** [5]: prior financial-domain language models that provide reference points for sentiment and financial-NLP tasks.
- **FinSim** [6]: semantic-similarity and term-classification tasks in the financial domain.

Earnings-call analysis, guidance extraction, and financial-document summarization are important downstream applications, but we treat them as deployment evaluations to be run on disclosed, held-out data with published metrics rather than as headline claims.

### 5.2 Reporting Protocol

For each benchmark we recommend disclosing the exact model version and quantization, the prompt template, the decoding parameters, the evaluation harness, and a side-by-side result for the unmodified Qwen3-8B base. For numerical tasks we recommend reporting tool-augmented and non-tool-augmented results separately, since the calculator tool—not the language model’s in-context arithmetic—is responsible for numeric correctness.

## 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 draft default-probability narratives, credit-grade rationales, and early-warning-signal summaries. These are intended to support, not replace, an institution’s existing credit models and underwriters; any deployment should validate the model’s outputs against the institution’s internal model on held-out portfolios before reliance. We report no accuracy or agreement figures in this report.

# 7 Operational Considerations

## 7.1 Latency Requirements

Trading and risk applications have strict latency requirements that differ by use case—from sub-second pre-trade compliance checks to multi-second portfolio risk narratives and longer research-report generation. As an 8B-parameter model, Zen Financial is small enough to serve interactively on a single modern accelerator, and quantized variants reduce latency further. We report no specific latency measurements in this report; achievable latency depends on the deployment hardware, quantization, batch size, and context length, and should be measured in the target environment.

## 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 an openly licensed general-purpose model can be adapted to financial language and workflows through targeted domain fine-tuning, while adding the tool-augmented quantitative reasoning, compliance-aware generation, and audit machinery that production financial services require. Built as an Apache-2.0 fine-tune of Qwen3-8B rather than a from-scratch or multi-hundred-billion-parameter system, it inherits a strong, transparently licensed foundation and remains small enough to deploy on commodity accelerators. We have intentionally avoided headline benchmark claims in this report; the honest measure of Zen Financial is a reproducible evaluation against public benchmarks such as FinQA and Financial PhraseBank, reported side-by-side with the unmodified Qwen3-8B base so that the contribution of the financial adaptation is transparent. Tool-augmented quantitative reasoning and compliance-aware generation address the key failure modes of general-purpose LLMs in financial applications.

## References

- [1] Qwen Team, Alibaba Group. Qwen3 Technical Report. 2025. Models released under the Apache-2.0 license. <https://github.com/QwenLM/Qwen3>
- [2] Chen, Z. et al. FinQA: A Dataset of Numerical Reasoning over Financial Data. *EMNLP*, 2021.
- [3] Malo, P. et al. Good Debt or Bad Debt: Detecting Semantic Orientations in Economic Texts. *JASIST*, 2014.
- [4] Araci, D. FinBERT: Financial Sentiment Analysis with Pre-Trained Language Models. *arXiv:1908.10063*, 2019.
- [5] Wu, S. et al. BloombergGPT: A Large Language Model for Finance. *arXiv:2303.17564*, 2023.
- [6] Xu, Y. et al. FinSim-4 Shared Task: Learning Semantic Similarities for the Financial Domain. *SIGIR-AP*, 2022.