CHAPTER 2

Prompt Engineering for Finance

كيف تكتب أوامر مالية فعّالة

Level: Beginner–Intermediate

Learning Objectives

- Master the DARE framework for structuring financial prompts
- Set effective context that shapes Claude's responses for finance-specific tasks
- Specify output formats that integrate directly into financial workflows
- Apply iterative refinement to improve model accuracy and completeness
- Avoid the most common prompt engineering mistakes in financial analysis

2.1 The DARE Framework

The DARE framework is a structured approach to writing financial prompts that consistently produce high-quality, actionable outputs. DARE stands for Define, Analyze, Request, Evaluate. Each component addresses a specific aspect of prompt construction that, when combined, dramatically improves the relevance and accuracy of Claude's responses.

The framework draws on established principles from financial analysis methodology and adapts them for AI interaction. Just as a well-structured investment memo follows a predictable format (situation, analysis, recommendation, risks), a well-structured financial prompt guides Claude toward a focused, professional output.

Component	Purpose	Key Questions
D — Define	Set the context, role, and scope	What industry? What company type? What time horizon? What currency? What accounting standard?
A — Analyze	Specify the data and analysis required	What data should Claude examine? What metrics matter? What comparisons are needed?
R — Request	State the exact output format and deliverables	What format? (table, JSON, memo) What sections? How detailed? What structure?
E — Evaluate	Instruct Claude to validate its own output	Does the BS balance? Are growth rates reasonable? Do ratios match industry norms? Any circular references?

2.2 D — Define: Setting the Context

The Define step establishes the analytical framework before any calculation begins. Without proper context, Claude defaults to generic responses. With precise context, Claude tailors its analysis to your specific situation.

Essential Context Elements

Role: Tell Claude what role to assume. "You are a senior equity research analyst" produces different output than "You are a credit analyst." The role shapes terminology, focus areas, and output style.

Industry: Specify the sector. A retail analysis requires different KPIs (same-store sales, GMROI) than a telecom analysis (ARPU, churn rate). Claude adjusts its framework accordingly.

Company type: Public vs. private, growth vs. mature, domestic vs. multinational. Each changes the relevant valuation approach and data availability.

Time horizon: Specify whether you need trailing twelve months (TTM), 5-year projections, or 10-year DCF. This determines the level of assumption detail required.

Currency and units: Always specify: "USD millions" or "SAR thousands." Ambiguity in units is one of the most common sources of error in AI-generated financial models.

Accounting standard: IFRS or US GAAP. This affects revenue recognition (IFRS 15 vs. ASC 606), lease accounting (IFRS 16 vs. ASC 842), and financial statement presentation.

Source: IFRS Foundation, ifrs.org; FASB Accounting Standards Codification, fasb.org

Chat — Complete Context Setting

```
You are a senior financial analyst at a Big Four advisory firm. I need
your help analyzing a mid-cap Saudi Arabian retail company (publicly
listed on Tadawul). Use:
- IFRS standards
- SAR (Saudi Riyals) in millions
- 5-year projection horizon (2025–2029)
- Benchmark against GCC retail sector averages
```

```
Confirm you understand this context before we proceed.
```

Expected Output: *Claude confirms the context parameters and may ask clarifying questions (e.g., sub-sector within retail, specific company name).*

Refinement: *If Claude does not confirm a parameter, restate it: "Please confirm: all figures in SAR millions, IFRS basis."*

API — System Prompt for Financial Context

```
system="You are a senior financial analyst specializing in GCC retail companies. Use IFRS standards. All currency in SAR millions. 5-year projection horizon. Benchmark against GCC retail sector averages from published sources. Always cite the source of any benchmark data you use. Output structured JSON with keys: context, assumptions, analysis, validation."
```

Expected Output: *The system prompt persists across all API calls in the session, ensuring consistent context.*

2.3 A — Analyze: Specifying the Data and Analysis

The Analyze step tells Claude exactly what data to examine and what analytical framework to apply. Vague instructions ("analyze this company") produce vague outputs. Specific instructions produce actionable analysis.

Structuring Your Analysis Request

There are three levels of analysis specificity, each appropriate for different situations:

Level 1: Exploratory: "Identify the key financial trends and potential concerns." Use when you want Claude to surface insights you may not have considered. Good for initial screening.

Level 2: Directed: "Analyze profitability trends (gross, operating, net margins) over 3 years and compare to sector median." Use when you know the area but want structured analysis.

Level 3: Precise: "Calculate ROE using 5-factor DuPont decomposition for FY2022–2024. Identify which factor drove the largest year-over-year change." Use when you need specific calculations with exact methodology.

For financial modeling, Level 2 and Level 3 prompts produce the most useful outputs. Level 1 is valuable for initial exploration but should be followed by more directed prompts.

Chat — Directed Analysis (Level 2)

```
Using the financial data I will provide, analyze:

1. Revenue growth trends (3-year CAGR and year-over-year)
2. Margin analysis (gross, EBITDA, operating, net) with trend direction
3. Working capital efficiency (DSO, DIO, DPO, cash conversion cycle)
4. Leverage metrics (D/E, net debt/EBITDA, interest coverage)
5. Return metrics (ROE, ROA, ROIC)

For each metric, indicate whether it is improving, stable, or
deteriorating. Flag any metric that deviates significantly from typical
retail benchmarks.
```

Expected Output: *A structured analysis with 5 sections, each containing calculated metrics, trend assessment, and benchmark comparison.*

Refinement: *Follow up: "Which single metric concerns you most, and why?"*

2.4 R — Request: Specifying Output Format

The Request step determines how Claude presents its analysis. The output format should match your downstream workflow. If you're feeding data into Excel, request a table. If you're building an automated pipeline, request JSON. If you're writing a memo, request structured prose.

Common Financial Output Formats

Structured Table: Best for: financial statements, ratio summaries, peer comparisons. Request: "Present as a table with columns [Year, Revenue, EBITDA, Net Income, FCF]."

JSON: Best for: API pipelines, database integration, programmatic processing. Request: "Return JSON with keys: assumptions{}, projections[], sensitivity_matrix[][]."

Executive Summary: Best for: board presentations, client memos. Request: "3-bullet executive summary followed by a 1-page detailed analysis."

Step-by-Step Calculation: Best for: learning, model documentation, audit trails. Request: "Show each calculation step with the formula, inputs, and result."

Waterfall / Bridge: Best for: variance analysis, revenue decomposition. Request: "Present as a waterfall: Prior Year → Price → Volume → Mix → FX → Current Year."

Chat — Table Format Request

```
Present a 5-year financial projection in the following table format:

| Line Item | FY2025E | FY2026E | FY2027E | FY2028E | FY2029E |
|-----------|---------|---------|---------|---------|---------|
| Revenue   |         |         |         |         |         |
| COGS      |         |         |         |         |         |
| Gross Profit |      |         |         |         |         |
| EBITDA    |         |         |         |         |         |
| Net Income|         |         |         |         |         |
| FCF       |         |         |         |         |         |
```

```
Below the table, list all assumptions used (growth rates, margins, tax
rate, etc.) with the source or basis for each assumption.
```

Expected Output: *A populated table with consistent formulas and a clearly labeled assumptions section.*

Refinement: *If any cell shows unusual growth or decline, ask: "Explain the year-over-year change in [line item] from FY20XX to FY20XX."*

API — JSON Format Request

```
User: "Project 5-year financials. Return JSON:
{
  "assumptions": {"revenue_growth": [], "gross_margin": [], "opex_pct":
[], "tax_rate": float},
  "projections": [{"year": int, "revenue": float, "cogs": float,
"gross_profit": float,
    "ebitda": float, "net_income": float, "fcf": float}],
  "validation": {"bs_balanced": bool, "growth_reasonable": bool, "notes":
[]}
}"
```

Expected Output: *Parseable JSON that can be directly loaded into Python (json.loads) or fed into an Excel automation pipeline.*

2.5 E — Evaluate: Validation and Cross-Checking

The Evaluate step is what separates professional financial modeling from casual analysis. By explicitly instructing Claude to validate its own output, you catch errors before they reach your model or client deliverable.

This step is inspired by the peer review process in investment banking, where models are reviewed by senior analysts before being presented. Claude can perform many of these checks automatically, but the human analyst must still verify the final output.

Essential Validation Checks

Balance sheet balance: Total assets must equal total liabilities plus equity. If Claude's projection doesn't balance, there is a structural error in the model.

Cash flow reconciliation: The change in cash on the balance sheet must equal net cash from the cash flow statement. This is a fundamental integrity check.

Growth rate reasonableness: Revenue growth should be compared to historical rates, industry averages, and GDP growth. A 50% growth rate for a mature utility company should be flagged.

Margin consistency: Margins should follow logical patterns. Gross margin should not increase if the company is lowering prices or facing input cost inflation without an explanation.

Circular reference detection: In linked financial models, interest expense depends on debt, which depends on cash flow, which depends on interest expense. Claude should identify and handle these correctly.

Sign conventions: Ensure consistency: are expenses shown as positive or negative? Are outflows in the cash flow statement positive or negative? Mixed conventions cause material errors.

Chat — Comprehensive Validation Request

```
Now validate the model you just built. Check:
```

```
1. Does the balance sheet balance in every projected year?
2. Does the change in cash reconcile between the BS and CF statement?
3. Are revenue growth rates reasonable vs. industry (compare to published
benchmarks)?
4. Do margins follow a logical trajectory? Flag any abrupt changes.
5. Is the interest expense calculation consistent with the average debt
balance?
6. Are there any circular references? If so, how did you resolve them?

Present your findings in a table: Check | Status (Pass/Fail/Warning) |
Details
```

Expected Output: *A validation table with pass/fail status for each check, plus explanatory notes for any warnings or failures.*

Refinement: *If any check fails, ask: "Fix [specific issue] and regenerate the affected statements."*

2.6 Common Mistakes and How to Fix Them

Mistake: Vague context

Bad: “Analyze this company.”

Good: “You are an equity research analyst. Analyze this Saudi retail company using IFRS, SAR millions, 5-year horizon.”

Why it matters: Without context, Claudeguesses the framework, which introduces inconsistency.

Good: “You are an equity research analyst. Analyze this Saudi retail company using IFRS, SAR millions, 5-year horizon.”

Why it matters: Without context,

Fix: Always start with the Define step. Specify role, industry, standards, currency, and time horizon.

Mistake: No output format

Bad: “Project the financials.”

Good: “Project 5-year financials in a table with columns: Year, Revenue, EBITDA, Net Income, FCF.”

Why it matters: Claudemay output prose when you need a table, or a table when you need JSON.

Good: “Project 5-year financials in a table with columns: Year, Revenue, EBITDA, Net Income, FCF.”

Why it matters:

Fix: Always include the Request step. Specify exact format, columns, and structure.

Mistake: Accepting first output

Problem: Claude’s first response may contain errors or suboptimal assumptions. Why it matters: Financial models must be reviewed and refined, whether built by a human or AI.

Fix: Always apply the Evaluate step. Request validation, then iterate on any issues found.

Mistake: Hardcoding assumptions

Bad: “Assume 10% growth.”

Good: “Use historical 3-year CAGR as the base case growth rate. Create bull (+2pp) and bear (–2pp) scenarios.”

Why it matters: Data-driven assumptions are more defensible and realistic.

Fix: Ground assumptions in data. Reference historical trends, industry benchmarks, or management guidance.

Mistake: Ignoring limitations

Problem: Treating Claude’s output as verified fact without cross-referencing. Why it matters: Claude may hallucinate benchmark data or make calculation errors.

Fix: Verify all numerical outputs independently. Cross-reference benchmarks with published sources (Damodaran, S&P, Bloomberg).

2.7 Prompt Templates for Common Financial Tasks

The following templates apply the DARE framework to the most common financial modeling tasks. Each template is ready to use—replace the bracketed placeholders with your specific data.

Template: Financial Statement Analysis

💬 Chat — Financial Statement Analysis

```
Define: You are a [Role] analyzing [Company] in the [Industry] sector.
Use [IFRS/GAAP], [Currency] in [Units].

Analyze: Review the income statement, balance sheet, and cash flow
statement for FY[Year]. Calculate: gross margin, EBITDA margin, net
margin, current ratio, quick ratio, D/E ratio, ROE, ROA, ROIC.

Request: Present results in a ratio summary table. Include a column for
industry benchmark comparison. Below the table, provide 3 key insights
and 1 risk factor.

Evaluate: Verify all ratios are mathematically correct. Flag any ratio
that deviates >20% from the industry benchmark.
```

Expected Output: *A complete ratio analysis table with benchmark comparison, insights, and risk assessment.*

Template: Revenue Forecasting

💬 Chat — Revenue Forecasting

```
Define: [Industry] company with [X] operating units. Historical revenue:
FY[Y1] $[A]M, FY[Y2] $[B]M, FY[Y3] $[C]M.

Analyze: Calculate historical CAGR. Identify growth drivers (price vs.
volume vs. mix). Assess market conditions and competitive position.

Request: Project revenue for 5 years using bottom-up methodology: [units]
× [revenue per unit]. Create 3 scenarios (Base/Bull/Bear) with clearly
```

```
labeled assumptions for each. Present in table format.

Evaluate: Verify that base case growth is within ±2% of historical CAGR. Confirm bull and bear cases represent reasonable upside/downside (not extreme outliers).
```

Expected Output: *Three scenario projections with assumptions, all grounded in historical data.*

Template: DCF Valuation

Chat — DCF Valuation

```
Define: Valuing [Company] in [Industry]. Use [IFRS/GAAP]. Currency: [X]. Shares outstanding: [Y]M.

Analyze: FCF Year 0 = $[A]M. Cost of equity (from CAPM): risk-free [Rf]%, beta [B], market premium [Mp]%. Cost of debt: [Kd]%. Tax rate: [T]%. Capital structure: [E]% equity, [D]% debt.

Request: Build a 5-year DCF model. Show: (1) WACC calculation, (2) FCF projections, (3) Terminal value (perpetuity growth method, g=[X]%), (4) PV of FCFs + TV, (5) Enterprise value, (6) Bridge to equity value, (7) Implied price per share. Format as step-by-step calculation.

Evaluate: Verify WACC is between 6–15% (flag if outside). Confirm terminal value is 50–80% of total enterprise value (flag if outside this range). Check that implied price is within ±30% of current market price.
```

Expected Output: *A complete DCF model with every calculation step shown and validated.*

Template: Peer Comparison

API — Peer Comparison

```
system="Senior equity analyst. Compare companies using standardized multiples. Cite data sources."

user="Compare these 5 [Industry] companies:
1. [Company A]: EV $[X]B, EBITDA $[Y]M, NI $[Z]M, Revenue $[W]M
2. [Company B]: ...
[...]
```

```
Calculate for each: EV/EBITDA, P/E, EV/Revenue, P/B.
Output JSON: {companies: [{name, ev_ebitda, pe, ev_rev, pb}], medians:
{ev_ebitda, pe, ev_rev, pb}, implied_valuation: {target_company,
range_low, range_mid, range_high}}"
```

Expected Output: *Structured JSON with peer multiples, sector medians, and implied valuation range.*

2.8 Iterative Refinement: The Conversation Approach

Financial modeling with Claude is rarely a single-prompt exercise. The most effective approach is a structured conversation where each prompt builds on the previous output. This mirrors the traditional analyst workflow: build, review, refine, validate.

The 4-Turn Financial Conversation Pattern

Turn 1: Context + Initial Build: Set the DARE context and request the initial model or analysis.

Turn 2: Review + Adjust Assumptions: "The revenue growth seems aggressive. Reduce base case to [X]% and show how it affects EBITDA and FCF."

Turn 3: Deepen + Extend: "Now add a sensitivity table showing how equity value changes across WACC (7–13%) and terminal growth (1–5%)."

Turn 4: Validate + Finalize: "Validate the complete model. Check all balances, growth reasonableness, and margin consistency. List any issues."

This pattern can extend to more turns for complex models. The key principle is that each turn has a clear purpose: build, adjust, extend, or validate. Avoid combining multiple objectives in a single prompt, as this increases the risk of errors.

Key Takeaways

- The DARE framework (Define, Analyze, Request, Evaluate) provides a repeatable structure for every financial prompt.
- Context setting (Define) is the single most impactful step. Specifying role, industry, standards, and currency eliminates ambiguity.
- Always specify your output format (Request). Match the format to your workflow: tables for Excel, JSON for automation, prose for memos.
- Never accept Claude's first output without validation (Evaluate). Request explicit cross-checks on balance, growth rates, and margin logic.
- Use iterative refinement: build in Turn 1, adjust in Turn 2, extend in Turn 3, validate in Turn 4.
- Ground all assumptions in data. Reference historical trends and published benchmarks rather than arbitrary numbers.
- The templates in Section 2.7 are ready to use. Replace bracketed placeholders with your specific data.

كيف تكتب أوامر مالية فعّالة

DARE Framework — إطار DARE

Define — حدّد

Analyze — حلّل

Request — اطلب

Evaluate — قيّم

Prompt Engineering — هندسة الأوامر

Context Setting — تحديد السياق

Output Format — تنسيق المخرجات

Iterative Refinement — التحسين التكراري

Cross-Validation — التحقق المتقاطع

Sensitivity Analysis — تحليل الحساسية

Scenario Analysis — تحليل السيناريوهات

Revenue Forecasting — توقع الإيرادات

DCF (Discounted Cash Flow) — التدفقات النقدية المخصومة

WACC (Weighted Average Cost of Capital) — المتوسط المرجح لتكلفة رأس المال

Peer Comparison — مقارنة النظراء

Executive Summary — ملخص تنفيذي

Assumptions — الافتراضات

Circular Reference — مرجع دائري

Hallucination (AI) — هلوسة (ذكاء اصطناعي)

References

IFRS Foundation. (2024). IFRS Standards Overview. ifrs.org

Financial Accounting Standards Board. (2024). ASC Codification. fasb.org

Anthropic. (2024). Claude API Documentation. docs.anthropic.com

Damodaran, A. (2024). “Valuation: Lecture Notes.” NYU Stern School of Business. pages.stern.nyu.edu/~adamodar

CFA Institute. (2024). “Global Investment Performance Standards.” cfainstitute.org