


FINANCIAL MODELLING GUIDE FOR CLAUDE

From a prompt to a controlled model using 23 specialist agents.

23 bounded agents | SET UP → BUILD → FORECAST → VALUE → STRESS-TEST → COMMUNICATE



Claude becomes useful in consequential work when prompts become governed work orders, repeatable methods become Skills, and narrowly scoped agents hand checked analysis to human reviewers.

ORIA

AI for complex slides - after the analysis is approved.

01 | OPERATING SYSTEM

Build the system, not a longer prompt.

A strong prompt improves one answer. A governed operating model improves the next hundred assignments - while keeping authority, evidence, versions, and review visible.

PROJECT CONTEXT	Mandate, definitions, policies, source hierarchy, templates, and disclosure rules that remain stable across runs.
WORK ORDER	This run's objective, evidence, method, deliverables, permissions, stop conditions, review gates, and acceptance tests.
SKILL	A versioned method for work you repeat: source ledger, model audit, scenario design, storyline, or deck QC.
AGENT	A bounded workflow owner with instructions, permitted tools, an output contract, and a named handoff - never a decision-maker.
CONNECTOR / MCP	Governed access to approved current data or systems. Start read-only; access is not proof that data is correct or entitled.

SET UP ONE AGENT BEFORE YOU BUILD THE LIBRARY

- Choose one recurring decision and define the artifact, evidence, tie-outs, reviewer, and prohibited actions.
- Give the agent the smallest tool and Skill surface it needs; stop on missing evidence, conflicts, or permission ambiguity.
- Test the standard case, a conflicting-input case, and a case that must escalate. Version the definition and record the owner.

OFFICIAL SOURCES

[Anthropic: effective agents](#)
[Claude Code: subagents](#)
[Agent Skills overview](#)

02 | PROMPTS + SKILLS

Write the work order. Record the repeatable method.

The best finance prompt is inspectable: it makes the decision, source boundary, method, authority, and definition of done explicit before Claude calculates or writes.

01 OBJECTIVE + DECISION

Question, owner, deadline, and action the work may support.

02 DELIVERABLES

Exact model, memo, storyboard, tables, schema, and file locations.

03 EVIDENCE

Approved sources, versions, as-of date, permissions, and conflict hierarchy.

04 METHOD + CONVENTIONS

Calculations, scenarios, units, signs, periods, accounting basis, and materiality.

05 AUTHORITY + GATES

What Claude may change, prohibited actions, stops, reviewers, and approvers.

06 ACCEPTANCE TESTS

Tie-outs, tolerances, reproducibility, citations, and cross-artifact consistency.

COPY-READY FIRST RESPONSE

Restate the objective, decision use, approved evidence, conventions, deliverables, permissions, stops, and acceptance tests. List missing inputs and conflicts. Propose the smallest agent route and review gates. Return the source ledger before analysis. Label material claims FACT, CALCULATION, ASSUMPTION, or HYPOTHESIS. Do not calculate, edit, publish, contact anyone, or take an external action until the named reviewer approves the route.

RECORD A SKILL WHEN THE METHOD SURVIVES REVIEW

- Capture why each source is trusted, how periods and units are normalized, where judgment enters, which exceptions stop the run, and what done means.
- Keep the stable procedure in SKILL.md; put detailed references, deterministic checks, and templates in references/, scripts/, and assets/ only when needed.
- Test normal, edge, stale-source, negative-trigger, and approval-required cases. Remove secrets and sensitive examples before recording or sharing.
- For slides, record Ori's audience, action-title, exhibit, sourcing, template, and deck-QC method - then pass only approved claims into production.

OFFICIAL SOURCES

[Prompting best practices](#)
[Skill authoring](#)
[Create custom Skills](#)

Use the smallest specialist set that can finish the work.

01-06 of 23 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

01

Name: MODEL SETUP

What it does: defines structure, interfaces, sequence, scenarios, controls, and gates.

When to use: new, extended, repurposed, or opaque models.

[Download the companion agent pack](#)

02

Name: EXCEL BEST PRACTICES

What it does: sets formula, input, sign, unit, link, and version conventions.

When to use: before Build or when standards are inconsistent.

[Download the companion agent pack](#)

03

Name: ACCOUNTING REFRESHER

What it does: creates the source ledger and maps accounting treatments into model lines.

When to use: when reporting changes or judgments affect comparability.

[Download the companion agent pack](#)

04

Name: THREE-STATEMENT MODEL

What it does: integrates approved schedules, statements, and cash/debt links.

When to use: after statement handoffs and approved refreshes.

[Download the companion agent pack](#)

05

Name: REVENUE DRIVERS

What it does: builds segment revenue from approved operating drivers.

When to use: when drivers can replace a top-line plug.

[Download the companion agent pack](#)

06

Name: OPERATING COSTS

What it does: builds COGS, headcount, opex, margins, and EBIT.

When to use: for fixed/variable or operating-leverage logic.

[Download the companion agent pack](#)

Use the smallest specialist set that can finish the work.

07-12 of 23 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

07

Name: WORKING CAPITAL

What it does: builds operating balances, drivers, and cash effects.

When to use: when cash conversion matters.

[Download the companion agent pack](#)

08

Name: CAPEX & D&A

What it does: builds capex, PP&E, D&A, and intangible roll-forwards.

When to use: when asset intensity affects results.

[Download the companion agent pack](#)

09

Name: DEBT SCHEDULE

What it does: builds debt, interest, liquidity, maturities, and covenants.

When to use: for financed or liquidity-sensitive models.

[Download the companion agent pack](#)

10

Name: TAX SCHEDULE

What it does: builds current, deferred, effective-rate, and cash tax.

When to use: when tax timing or attributes matter.

[Download the companion agent pack](#)

11

Name: INCOME STATEMENT

What it does: assembles earnings and diluted-share/EPS outputs.

When to use: after operating, debt, and tax handoffs.

[Download the companion agent pack](#)

12

Name: BALANCE SHEET

What it does: assembles assets, liabilities, equity, and retained earnings.

When to use: after schedules and opening balances reconcile.

[Download the companion agent pack](#)

05 | AGENT DIRECTORY | BUILD + FORECAST + VALUE

Use the smallest specialist set that can finish the work.

13-18 of 23 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

13

Name: CASH FLOW STATEMENT

What it does: assembles cash flows and reconciles ending cash.

When to use: after statement movements are traceable.

[Download the companion agent pack](#)

14

Name: BUDGETING & FORECASTING

What it does: calibrates drivers and controls sourced, owned, versioned assumptions.

When to use: for base forecasts, budgets, updates, and cases.

[Download the companion agent pack](#)

15

Name: DCF VALUATION

What it does: builds discounting, terminal value, and EV-to-equity bridge.

When to use: with audited cash flow and approved WACC/terminal method.

[Download the companion agent pack](#)

16

Name: COMPARABLE COMPANIES

What it does: normalizes peers and applies consistent trading multiples.

When to use: with approved peers and current sourced data.

[Download the companion agent pack](#)

17

Name: M&A MODEL

What it does: structures sources/uses, purchase accounting, synergies, and accretion/dilution.

When to use: for an authorized acquisition case.

[Download the companion agent pack](#)

18

Name: LBO MODEL

What it does: structures entry/exit and sponsor MOIC/IRR around agent 09 debt.

When to use: for authorized leveraged-return analysis.

[Download the companion agent pack](#)

Use the smallest specialist set that can finish the work.

19-23 of 23 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

19

Name: MODEL CIRCULARITY

What it does: implements approved circularity and tests convergence/switches.

When to use: for genuine cash/debt/interest feedback loops.

[Download the companion agent pack](#)

20

Name: SCENARIO ANALYSIS

What it does: turns risks into driver-based cases and reverse stresses.

When to use: for resilience, liquidity, covenants, or thesis risk.

[Download the companion agent pack](#)

21

Name: SENSITIVITY ANALYSIS

What it does: builds sensitivities and tests base ties/direction.

When to use: when uncertain inputs drive outcomes.

[Download the companion agent pack](#)

22

Name: MODEL CHECKS

What it does: audits formulas, links, provenance, controls, and end state.

When to use: after integration/changes and before communication.

[Download the companion agent pack](#)

23

Name: OUTPUTS & AUTOMATION

What it does: packages approved outputs and an Oria-ready decision brief.

When to use: only after qualified Gate 3 approval.

[Download the companion agent pack](#)

Run the first controlled sprint with a minimum viable team.

The coordinator owns versions, dependencies, approvals, and open items. Specialist agents receive frozen inputs, return structured artifacts, and stop at named human gates.

01	BRIEF	Lock the decision, source boundary, output contract, permissions, and tests.
02	EVIDENCE	Build the source ledger first; resolve conflicts before synthesis.
03	BUILD	Route only independent tasks in parallel; preserve versions and handoffs.
04	CHALLENGE	Use a separate reviewer to reproduce material outputs and attack assumptions.
05	APPROVE	A named professional approves numbers, narrative, confidentiality, and release.
06	ORIA	Hand approved page briefs to Oria for complex, editable slides; then run deck QA.

30-DAY CONTROLLED PILOT

WEEK 1	Select one recurring workflow; freeze sources, baseline quality, reviewers, and failure tests.	WEEK 2	Configure the smallest agent route; keep tools read-only; test conflicts and stops.
WEEK 3	Run live work beside the current process; log corrections, time, provenance, and escalations.	WEEK 4	Approve, revise, or retire the route. Scale only after control and decision quality improve.

THE ORIA PRODUCTION BRIEF

Audience + decision → action-title spine → one message per slide → decisive exhibit + exact source IDs → caveat + review status → appendix proof. Claude structures the approved argument; Oria turns it into beautiful, complex slides.

Get the companion agent pack

Build complex slides with Oria

OFFICIAL SOURCES

- Managed-agent permissions
- Multiagent orchestration
- Financial services repository