PROPERTY ANALYSIS WITH CLAUDE

A practitioner guide to Claude for CRE underwriting tasks

Better context. Better conviction.

Sample — Executive Summary & Section 1

Forthcoming Q3 2026 · Free · Editorial board: Mark Kuklis

Property Analysis with Claude

Model specs, context windows, and pricing current as of May 2026. Verify before deploying to a production workflow: platform.claude.com/docs/en/about-claude/models/overview

Language models have moved faster than most CRE practitioners' mental models of them. Claude, developed by Anthropic, is now capable enough that a senior analyst can reduce the labor cost of document-intensive workflows with it -- rent roll parsing, lease abstraction, closing-package Q&A -- and precise enough in structured-output mode that results land in a usable format rather than cosmetically formatted noise. It is also capable of hallucinating a lease option that does not exist, confidently citing a section number that is wrong, and producing a summary that invents detail not present in the source document.

This brief covers both sides.

EXECUTIVE SUMMARY

Honest verdict: Claude is the stronger default for CRE document analysis. Its context window accommodates large closing packages in a single call; its structured-output compliance on complex nested schemas is reliable; its hallucination pattern tends toward omission rather than confident fabrication, which makes errors easier to catch.

What it does not do: live market data access, spreadsheet computation without a separate tool, or legal interpretation. The 95% field-accuracy figure on standard rent rolls reflects clean-format inputs -- heavily negotiated leases with exhibit-referenced options are a different story.

Decision rule: Use Claude (Sonnet 4.6 or Opus 4.x -- *current as of May 2026; verify at platform.claude.com*) for long-document analysis, structured extraction from complex lease documents, closing-package Q&A, and investor memo drafting. Use ChatGPT where you need numerical computation on uploaded spreadsheets. See the companion brief and the task matrix below for the full cross-platform decision rule.

Non-negotiable: human review before any output enters an underwriting model or a client deliverable. Neither model is a terminal step.

This brief is organized around the eight workflows where Claude delivers the most value in CRE underwriting and transaction work, the prompt patterns that improve reliability in each, and the tasks that remain unsafe to delegate at current capability levels. It is not a product endorsement.

1. The model family -- what you are actually using

Anthropic's current Claude family as of May 2026 consists of three tiers, all operating under the Claude 4.x generation.

Opus 4.x is the highest-capability tier. Current Opus models carry a 1 million-token context window on the Claude API, Amazon Bedrock, and Vertex AI, with a maximum output of 128,000 tokens. This is the model to reach for when running a full loan agreement or a 200-page OM through a single call. Pricing as of May 2026 is approximately $5 per million input tokens and $25 per million output tokens -- *verify at platform.claude.com/docs/en/about-claude/pricing*.

Sonnet 4.6 *(current as of May 2026 -- verify before use)* is the production workhorse. It carries a 1 million-token context window on the API and most cloud platforms, with 64,000-token maximum output. Pricing is approximately $3 per million input tokens and $15 per million output tokens. For the majority of document analysis workflows -- lease abstraction, rent roll parsing, OM Q&A -- Sonnet 4.6 is the appropriate default; its quality is close to Opus at meaningfully lower cost and faster latency.

Haiku 4.5 *(current as of May 2026 -- verify before use)* is the lightweight tier, with a 200,000-token context window, priced at approximately $1 input / $5 output per million tokens. It is appropriate for high-volume simple tasks: extracting a date range from a short lease exhibit, classifying document types at intake, pre-screening a stack of files. For substantive

analysis it is undersized.

One operational note: the 1 million-token Sonnet 4.5 beta was retired on April 30, 2026. If you have workflows built around `claude-sonnet-4-5` with extended context, verify your model string before using a long-document workflow at scale.

Access is through three channels: Claude.ai (web interface, personal and team plans), the Anthropic API (direct programmatic access, requires developer setup), and enterprise deployments via Amazon Bedrock or Google Vertex AI. For most CRE practitioners, Claude.ai Pro or Team is the practical entry point; API access becomes worthwhile when you are running the same extraction task against hundreds of documents.

That’s the sample.

The full **Property Analysis with Claude** brief runs the complete reference end to end — every section, all tables, full citations.

Forthcoming Q3 2026 · Free

Join the waitlist at valoreregistry.com to be notified at release.