Alphabet: Search Economics in the Age of Gemini

Defensive Innovation or Self-Cannibalization—and What It Means for Valuation

Case Packet (Student Handout)

Course: AI for Business / Corporate Valuation with LLMs

Prepared by: Miao Liu

Date: March 31, 2026

Teaching-use only

- **No proprietary data required.** Exhibits are assembled from public sources (investor materials, EDGAR filings, and official product documentation/blog posts).
- **LLM use is integral to the learning goals** and will be assessed via (i) evidence discipline and (ii) reproducibility (prompt audit trail).
- **Golden rule: LLMs are not data sources.** All numbers must be taken from exhibits and/or computed in your spreadsheet.

Student Handout

0.1 Case Overview

Alphabet's core economic engine has historically been **Google Search monetization**: user queries drive ad impressions and clicks, which generate advertising revenue. In the last two years, generative AI has begun to change *how* search is experienced through features such as **AI Overviews** and **AI Mode**, increasingly powered by the **Gemini** model family.

This shift creates a valuation puzzle:

- If AI answers make search *more useful*, query volume and engagement could rise.
- If AI answers reduce outbound clicks, ad load and click-through rates could fall—a form of **self-cannibalization**.
- AI answers are expensive to serve. Alphabet is scaling technical infrastructure for AI, implying **large capital expenditures and depreciation** that affect near-term free cash flow and long-run reinvestment needs.
- At the same time, Google Cloud may benefit from enterprise AI infrastructure and solutions demand.

In February 2026, Alphabet reported Q4 and full-year 2025 results and disclosed an unusually explicit **CapEx investment range for 2026** (**Exhibit** 1; **Exhibit** 2). You are asked to translate this AI-driven product transition into a valuation narrative that is *mechanistic* (drivers → unit economics → margins → free cash flow) rather than purely qualitative.

Decision Prompt (the "IC question")

If AI changes how people search, is Alphabet executing **defensive innovation** that protects and expands Search economics—or triggering **self-cannibalization** that compresses margins and increases reinvestment needs?
Does AI raise or lower Google's terminal value?

0.2 Role and Decision Context

Assume you are an equity research associate at a long-only fund. Your portfolio manager holds Alphabet and asks you to produce a short investment-committee-style write-up that answers:

- What are the key mechanisms by which AI could change Search economics?
- How does the AI-driven CapEx profile affect the *timing* of free cash flow?
- Is Cloud's AI-driven growth sufficient to offset Search monetization risk (or amplify it via infrastructure spending)?
- What are the **2–3 most valuation-sensitive assumptions** right now?

0.3 Learning Goals

By the end of the case, you should be able to:

- Translate an LLM-enabled product shift into a **driver tree** that connects user behavior to monetization, cost, margins, and free cash flow.
- Use exhibits to build an **evidence-based fact bank** (facts vs. assumptions vs. inferences).
- Design a **sensitivity grid** for high-uncertainty valuation: (CapEx intensity × Search margin impact × Cloud growth).
- Produce an **auditable valuation memo section** with explicit assumptions, caveats, and cross-checks.
- Use an LLM as a **research and structuring co-pilot** while maintaining an audit trail and avoiding fabricated data.

0.4 What You Are Given (Case Materials)

You are provided:

- A set of public exhibits (Section 0.7) including Alphabet's earnings release and selected product documents on AI in Search.
- An Excel template (or you may build your own) with suggested tabs:
  - **Timeline** (Q1)
  - **Economics Snapshot** (Q2)
  - **Fact Bank** (Q3)
  - **Driver Tree** (Q4)
  - **Sensitivity Grid** (Q5)
  - **Memo Draft** (Q6)
  - **Prompt Log** (required)

LLM Workflow Rules (Required)

- **No fabricated numbers.** If a number is not in the exhibits (or in your spreadsheet), do not use it.
- **Separate** (FACT) vs. (ASSUMPTION) vs. (INFERENCE). Every numeric input in your spreadsheet must be tagged.
- **Prompt audit trail.** Save prompts used, key excerpts of model outputs, and verification notes (what you checked or corrected).
- **Model role discipline.** Use the LLM for: extracting statements, summarizing disclosures, structuring driver trees, designing sensitivity plans, and critiquing your memo—not for "finding" market data.

0.5 Case Tasks and Questions

Q1 — Build the high-signal AI → economics timeline (LLM-assisted, human-verified)

Use the **Timeline** tab in the spreadsheet.

Using Exhibits 1–7, build a **5–8 row timeline** covering the period from **May 2024 to February 2026**. For each row, include:

- date,
- event (product, monetization, infrastructure, or disclosure),
- why the event matters analytically,
- one exact quote or one exact number,
- primary citation (**Exhibit** # page, timestamp, or line range).

Your timeline should tell a coherent story: the evolution of AI in Search (AI Overviews / AI Mode / Gemini), how Alphabet frames monetization, and how infrastructure investment (CapEx, depreciation) scales to support AI.

Q2 — Spreadsheet A: quantify baseline exposure to Search, Cloud, and CapEx

Use the **Economics Snapshot** tab.

Using **Exhibit** 1 and **Exhibit** 2, compute the following **four** baseline metrics (you may add others if useful):

1. **Search growth (Q4 YoY):** Google Search & other revenue growth from Q4 2024 to Q4 2025.
2. **Cloud margin (Q4):** Google Cloud operating margin (operating income / revenue) for Q4 2025.
3. **TAC intensity (Q4):** total traffic acquisition costs (TAC) as a percentage of Google advertising revenue (and briefly interpret what could happen to TAC under AI Mode).
4. **CapEx shock (2026 vs. 2025):** compare 2025 CapEx (cash purchases of property and equipment) to the 2026 CapEx investment range; compute the implied increase in dollars and as a percentage of 2025 revenues.

Interpretation prompt: In one paragraph, explain why each metric is relevant to the valuation question (terminal value vs. near-term FCF timing).

Q3 — Build the Fact Bank (the "no-hallucination" foundation)

Use the **Fact Bank** tab.

Using **Exhibit** 1–**Exhibit** 3, extract **12–15** factual statements that are decision-relevant for valuing Alphabet in the age of Gemini. Organize them into four buckets:

1. **Search product shift (AI Overviews / AI Mode / Gemini integration)**
2. **Search monetization and pricing language** (what management says about monetization sequencing, ad formats, or "monetize differently")
3. **Infrastructure + CapEx + depreciation** (what is scaling and why)
4. **Cloud + enterprise AI demand** (growth, backlog, solutions, margins)

Each row must include: (i) the statement, (ii) label (FACT / ASSUMPTION / INFERENCE), (iii) where it appears (**Exhibit** #), and (iv) why it matters.

LLM prompt scaffold for Q3 (Fact Bank)

Role: "You are a financial analyst assistant. Use only the text I provide."
Task: "Extract 12–15 valuation-relevant factual statements about AI in Search, monetization, Cloud AI demand, and CapEx guidance. Output a table: [Statement] [Category] [Why it matters] [Exact citation placeholder]. Do *not* invent numbers or facts. If the text is ambiguous, mark it as INFERENCE and explain why."
Note: You must still verify each statement against the exhibit and fill in the citation yourself.

Q4 — Build the Driver Tree: AI → engagement → monetization → margin → FCF

Use the **Driver Tree** tab.

Create a driver tree that connects the AI search product shift to valuation. Your tree must:

- Start with **AI product features** (AI Overviews, AI Mode, Gemini capabilities).
- Map to **user behavior metrics** (query volume, query complexity, session length, retention).
- Map to **monetization mechanics** (ad load, CTR, CPC, conversion, TAC rate, new ad formats).
- Map to **unit cost and margin** (compute per query/inference; depreciation; opex leverage).
- End with **free cash flow** (operating cash flow, CapEx, working capital timing).

Minimum requirement: At least **3 layers** deep and at least **10 nodes** total. For each edge, write a one-line economic rationale and propose **one measurable proxy**.

Hint: "Monetization follows UX" is not a valuation model

Management may emphasize improving user experience before monetization. As an analyst, your job is to translate that sequencing into *explicit* assumptions: what changes *when*, and what is the impact on margins and reinvestment?

Q5 — Design the Sensitivity Grid (3D uncertainty)

Use the **Sensitivity Grid** tab.

Build a scenario/sensitivity design with three dimensions:

1. **CapEx intensity:** low / base / high within the disclosed 2026 CapEx investment range.
2. **Search margin impact:** a range of Search (or Services) margin outcomes reflecting AI-driven cannibalization vs. improved monetization (e.g., 0 bps, −100 bps, −200 bps vs. your baseline proxy).
3. **Cloud growth:** a range of Cloud revenue growth outcomes (e.g., baseline growth proxy from Q4 plus upside/downside cases).

Deliverable: A table that defines each scenario, the exact assumption changes, and what output you will track (e.g., 2026 FCF, FCF margin, implied EV using a simple multiple, or a one-year DCF bridge).

LLM prompt scaffold for Q5 (Sensitivity design + Excel implementation)

Role: "You are a valuation analyst and Excel modeling coach."
Inputs: "Here are my baseline metrics and my Driver Tree."
Task: "Propose a 3D sensitivity grid on (CapEx intensity, Search margin impact, Cloud growth). For each variable, suggest a justified range anchored to my baseline. Then provide an Excel-ready structure: named inputs, scenario table layout, and formulas to compute a chosen output (e.g., FCF or implied EV). Do not introduce new data; use placeholders if needed and label them ASSUMPTION."
Constraint: "Keep it implementable in under 30 minutes."

Q6 — Write the Investment Committee memo section (terminal value call)

Write a **400–600 word** memo section titled:

"Alphabet: Does AI raise or lower Google's terminal value?"

Requirements:

- Include a clear thesis (1–2 sentences) and a recommended stance (e.g., maintain multiple, compress multiple, widen valuation range).
- Cite at least **4 facts** from your Fact Bank (**Exhibit**-based).
- Cite at least **2 spreadsheet metrics** from Q2 and/or scenario outputs from Q5.
- Explicitly list your **top 3 valuation-sensitive assumptions** and how you would monitor them over the next 2–4 quarters.
- Include a short "Risks and Caveats" paragraph that distinguishes **execution risk**, **monetization risk**, and **infrastructure/CapEx risk**.

Optional extension (for stronger submissions)

Add a short "valuation method note" explaining whether you would rely more on a DCF (CapEx and reinvestment explicit) or on multiples (peer comps) in this setting—and why.

0.6 Required Deliverables

Deliverable	Format / length	What it must include
Report	Max 3 pages (appendices excluded)	Clear responses to Q4–Q6, with explicit assumptions and exhibit citations
Completed spreadsheet	Excel spreadsheet using the provided template (or equivalent)	Completed **Timeline**, **Economics Snapshot**, **Fact Bank**, **Driver Tree**, and **Sensitivity Grid** tabs with clearly labeled exhibit inputs
LLM audit trail (required)	Appendix	prompts used + selected outputs + verification notes showing what you checked, rejected, or corrected

0.7 Exhibit Table of Contents

You are provided the following exhibits.

Exhibit	Description and source
Exhibit 1	Alphabet Announces Fourth Quarter and Fiscal Year 2025 Results (earnings release; includes Search, Cloud, TAC, and 2026 CapEx range). Source: Alphabet Investor Relations / EDGAR. https://www.sec.gov/Archives/edgar/data/1652044/000165204426000012/googexhibit991q42025.htm
Exhibit 2	Alphabet Q4 2025 Earnings Call (prepared remarks and Q&A transcript; includes CapEx, cash flow, Cloud backlog, and AI commentary). Source: Alphabet Investor Relations event page. https://abc.xyz/investor/events/event-details/2026/2025-Q4-Earnings-Call-2026-Dr_C033hS6/default.aspx
Exhibit 3	Alphabet Annual Report on Form 10-K for fiscal year 2025 (risk factors and MD&A; includes discussion of AI Overviews/AI Mode monetization and technical infrastructure investment). Source: Alphabet Investor Relations / EDGAR. https://s206.q4cdn.com/479360582/files/doc_financials/2025/q4/GOOG-10-K-2025.pdf
Exhibit 4	"Generative AI in Search: Let Google do the searching for you" (AI Overviews rollout and product framing). Source: Google Blog (Search). https://blog.google/products-and-platforms/products/search/generative-ai-google-search-may-2024/
Exhibit 5	"AI in Search: Going beyond information to intelligence" (AI Mode framing and usage). Source: Google Blog (Search). https://blog.google/products-and-platforms/products/search/google-search-ai-mode-update/
Exhibit 6	"Google brings Gemini 3 AI model to Search and AI Mode" (Gemini 3 integration and product details). Source: Google Blog (Search). https://blog.google/products-and-platforms/products/search/gemini-3-search-ai-mode/
Exhibit 7	"Just ask anything: a seamless new Search experience" (updates to AI Overviews and AI Mode; framing on user experience and follow-up queries). Source: Google Blog (Search). https://blog.google/products-and-platforms/products/search/ai-mode-ai-overviews-updates/

Optional background reading (not required, not graded)

Google Search Central documentation: "AI features and your website" (how AI Overviews and AI Mode present information and links).
https://developers.google.com/search/docs/appearance/ai-features