Tribunal Session Summary

Session: conclave-20260603-154424-provide-a-best-in-9c37

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $6.1695 | Time: 39m 44s

THE PROMPT: Provide a best in class estimate of the EBITDA and EBITDA margins for the following companies, decoposed by top line revenue, operarting expenses with tech and headcount broken out, depreciation, and other expenses, noting that they are private companies so you will need to do a lot of deep dives and research. - Semi Analysis (I have been "told" they do in excess of $30M in revenue) - Fund A AI - AlphaSense - Hebbia - Rogo - LinqAlpha - Fabricated Knowledge

THE RULING: The Tribunal ruled to **SYNTHESIZE** the positions of Claude Opus 4.7 and Claude Opus 4.7dvocate-B, producing unaudited, triangulated EBITDA estimates for companies with disclosed secondary data (SemiAnalysis, AlphaSense, Hebbia, Rogo, LinqAlpha) while declining estimates for entities without verifiable anchors (Fund A AI, Fabricated Knowledge). This synthesis reconciles the strict data constraints with the user's request for deep-dive analysis by leveraging disclosed public press releases, founder statements, and industry benchmarks to build scenario-based financial models. The decision prioritized decision-utility over strict epistemic purity, acknowledging the Current Facts' prohibition on filed financials but interpreting "deep dives" as authorization to use secondary sources with e...

The Prompt

Provide a best in class estimate of the EBITDA and EBITDA margins for the following companies, decoposed by top line revenue, operarting expenses with tech and headcount broken out, depreciation, and other expenses, noting that they are private companies so you will need to do a lot of deep dives and research. - Semi Analysis (I have been "told" they do in excess of $30M in revenue) - Fund A AI - AlphaSense - Hebbia - Rogo - LinqAlpha - Fabricated Knowledge

The Ruling

The Tribunal ruled to **SYNTHESIZE** the positions of Claude Opus 4.7 and Claude Opus 4.7dvocate-B, producing unaudited, triangulated EBITDA estimates for companies with disclosed secondary data (SemiAnalysis, AlphaSense, Hebbia, Rogo, LinqAlpha) while declining estimates for entities without verifiable anchors (Fund A AI, Fabricated Knowledge). This synthesis reconciles the strict data constraints with the user's request for deep-dive analysis by leveraging disclosed public press releases, founder statements, and industry benchmarks to build scenario-based financial models. The decision prioritized decision-utility over strict epistemic purity, acknowledging the Current Facts' prohibition on filed financials but interpreting "deep dives" as authorization to use secondary sources with explicit confidence labels.

The synthesis adopts Claude Opus 4.7's "triangulation" method—combining disclosed top-line data (ARR, funding), headcount data (LinkedIn), and industry cost benchmarks (loaded compensation, compute costs)—while using GPT-5's "per-$100 revenue archetypes" as a fallback for entities lacking specific anchors. All outputs are labeled "unaudited triangulated scenarios" to reflect the absence of GAAP filings. The resulting analysis identifies SemiAnalysis as a high-margin research boutique (35-50% EBITDA), AlphaSense as a heavy investment-stage entity burning cash due to headcount burden (-30% to -8% EBITDA), and the AI-native startups (Hebbia, Rogo, LinqAlpha) as growth-stage burners with deeply negative margins. Fund A AI was deemed unidentifiable, and Fabricated Knowledge was modeled as a micro-SaaS entity with owner-earnings rather than institutional EBITDA.

Core Analysis

Methodology: Triangulation and Archetypes

The synthesis combines two distinct analytical frameworks to address the data void. For companies with disclosed secondary data, it applies **Claude Opus 4.7's triangulation methodology**: anchoring revenue to public disclosures (press releases, founder statements), estimating personnel costs via headcount (LinkedIn) multiplied by industry-standard loaded compensation benchmarks (Radford, Levels.fyi), and modeling tech/compute costs based on disclosed spend or inference-heavy usage patterns. For entities lacking these anchors, it applies **GPT-5's per-$100 revenue archetypes**, using generic SaaS and research-boutique cost structures to bound plausible outcomes. All figures are explicitly labeled as "unaudited triangulated

scenarios" to signal the absence of primary financial statements.

Company-Specific Estimates

SemiAnalysis (The Research-IP Boutique)

- **Revenue**: $60–80M. This range anchors to the user's ">$30M" input and incorporates secondary reporting of a growth trajectory from ~$20M to a projected $100M, split at a conservative midpoint.
- **Operating Expenses**:
  - **Headcount**: Estimated at ~50 FTEs. Using a loaded cost of $200k–$250k (premium for SF-based analysts/engineers), personnel costs are ~$10–$12.5M.
  - **Tech & Compute**: ~$7M. This is a founder-disclosed figure (Dylan Patel) representing a significant jump in AI infrastructure spend, acting as a structural drag on gross margin.
  - **G&A/S&M**: Estimated at $13–$25M to support platform maintenance and growth, though distribution is largely Substack-led (low CAC).
- **Depreciation & Amortization**: Minimal (<$1M), primarily server hardware.
- **EBITDA**: $20–35M.
- **EBITDA Margin**: 35–50%.
- **Mechanism**: SemiAnalysis operates as a high-leverage intellectual property business. The massive revenue-per-employee ratio (~$1.2-1.6M/FTE) and low customer acquisition cost (CAC) drive software-like margins, distinguishing it from the heavy-burn AI SaaS cohort.

AlphaSense (The Heavy Investment Phase)

- **Revenue**: ~$500M ARR. Verified via company press release.
- **Operating Expenses**:
  - **Headcount**: 2,938 FTEs. Using a blended loaded comp benchmark of $180k–$200k (standard for enterprise SaaS with a mix of US/Global talent), personnel costs total ~$530–$590M.
  - **Tech & Compute**: Estimated at $20–$40M to support search and AI processing over a massive document corpus.
  - **Other OpEx**: S&M and G&A are largely embedded in the headcount figure for a firm of this size, with additional T&E and office overhead.
- **Depreciation & Amortization**: $10–$20M (including Tegus intangibles).
- **EBITDA**: –$150M to –$40M.
- **EBITDA Margin**: –30% to –8% (GAAP basis).
- **Mechanism**: AlphaSense is in a post-scale investment phase. The headcount cost burden alone approximates total revenue, structurally suppressing EBITDA. The $930M Tegus acquisition adds integration costs and data licensing COGS, further pressuring margins despite the massive top-line.

Hebbia (The Growth-Stage Burner)

- **Revenue**: $30–50M (Run-rate). Extrapolated from a disclosed $13M ARR (mid-2024) and subsequent funding/acquisition activity (FlashDocs).
- **Operating Expenses**:
  - **Headcount**: 153 FTEs. Loaded comp estimated at $250k–$300k (high density of AI engineering talent) = ~$38–$45M.

  - **Tech & Compute**: High inference costs for neural search estimated at $5–$10M.
  - **S&M**: Aggressive land-grab spend estimated at $2–$10M.
- **Depreciation & Amortization**: $1–$2M.
- **EBITDA**: –$15M to –$5M.
- **EBITDA Margin**: –30% to –15%.
- **Mechanism**: Hebbia exhibits classic growth-stage economics. Revenue is scaling, but the premium cost of AI talent and the "Inference Tax" (continuous LLM costs) keep margins deeply negative.

Rogo (The Land Grab)

- **Revenue**: $25–45M. Inferred from Series C stage ($160M raised) and competitive positioning in AI for finance.
- **Operating Expenses**:
  - **Headcount**: 156 FTEs. Loaded comp ~$250k+ = ~$39

How We Got Here

Dissenting Opinions

No dissenting opinions were filed in this session.