Tribunal Session Summary

Session: conclave-20260603-224307-provide-a-best-in-e7c1

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $8.8261 | Time: 40m 51s

THE PROMPT: Provide a best in class estimate of the EBITDA and EBITDA margins for each of the following companies, decomposed by top line revenue, operarting expenses with technology expense, headcount, depreciation, and other expenses broken out, noting that they are private companies so you will need to do a lot of deep dives and research. SemiAnalysis (I have been "told" they do in excess of $30M in revenue), Fund A AI (https://funda.ai/), AlphaSense, Hebbia, Rogo, LinqAlpha, and Fabricated Knowledge. Please do this for all seven of these companies. Note any similar companies that I am missing.

THE RULING: ACCEPT Claude Opus 4.7 (with specific reliance on GPT-5 for accounting methodology). The seven companies bifurcate sharply into two EBITDA profiles: **boutique research/content businesses** (SemiAnalysis, Fabricated Knowledge) generate high EBITDA margins (40–60%) due to asset-light models and elite IP leverage, while **venture-backed AI SaaS platforms** (AlphaSense, Hebbia, Rogo, LinqAlpha, Fund A AI) operate at significant losses (-15% to -150%) driven by aggressive headcount scaling, high inference costs, and growth prioritization.

The Prompt

Provide a best in class estimate of the EBITDA and EBITDA margins for each of the following companies, decomposed by top line revenue, operarting expenses with technology expense, headcount, depreciation, and other expenses broken out, noting that they are private companies so you will need to do a lot of deep dives and research. SemiAnalysis (I have been "told" they do in excess of $30M in revenue), Fund A AI (https://funda.ai/), AlphaSense, Hebbia, Rogo, LinqAlpha, and Fabricated Knowledge. Please do this for all seven of these companies. Note any similar companies that I am missing.

The Ruling

ACCEPT Claude Opus 4.7 (with specific reliance on GPT-5 for accounting methodology).

The seven companies bifurcate sharply into two EBITDA profiles: **boutique research/content businesses** (SemiAnalysis, Fabricated Knowledge) generate high EBITDA margins (40–60%) due to asset-light models and elite IP leverage, while **venture-backed AI SaaS platforms** (AlphaSense, Hebbia, Rogo, LinqAlpha, Fund A AI) operate at significant losses (-15% to -150%) driven by aggressive headcount scaling, high inference costs, and growth prioritization.

Why this answer: Claude Opus 4.7 provided the most rigorous evidence base, successfully defending the existence of the $930M Tegus acquisition with 10 primary sources and anchoring AlphaSense's revenue to verified public disclosures ($600M ARR). However, Claude Opus 4.7 initially committed a critical accounting error by including Depreciation & Amortization (D&A) in their EBITDA calculation. GPT-5 identified this error, forcing a correction that shifted AlphaSense's EBITDA from a deeply negative GAAP view to a more accurate Adjusted EBITDA range (-15% to -38%). The final "best in class" estimate is a synthesis of Claude Opus 4.7's research depth (Tegus verification, ARR anchors) and GPT-5's accounting rigor (EBITDA definition exclusion of D&A), supported by the qualitative bifurcation thesis introduced by Gemini 3.1 Pro.

Key Caveat: The Fresh Eyes review correctly noted that the Majority Opinion output was truncated and missing the requested "similar companies" section. While the judicial verdict on the advocate's position stands, the final deliverable failed to fully satisfy the formatting requirements of the original prompt.

Core Analysis

The Bifurcation Thesis

The deliberation established a structural divide in the AI research vertical. **Bootstrapped, analyst-driven platforms** (SemiAnalysis, Fabricated Knowledge) operate with near-zero marginal cost of distribution, allowing them to retain 40–60% of revenue as EBITDA once they scale past ~$20M. In contrast, **VC-backed enterprise AI SaaS** (AlphaSense, Hebbia, Rogo) are trapped in negative EBITDA territory because their business models require heavy "forward-deployed" engineering teams to implement bespoke workflows and

massive spend on data licensing and inference compute. The decisive variable is not revenue scale, but the ratio of headcount cost to revenue.

Company-by-Company Decomposition

1. SemiAnalysis

- **Revenue:** $25M – $35M (FY25), projecting toward $100M in 2026.
- **EBITDA:** +$10M – +$15M.
- **EBITDA Margin:** 40% – 55%.
- **Decomposition:**
  - **Headcount:** ~50 FTEs (corrected from initial 15–20 estimates based on LinkedIn/video evidence). Fully-loaded comp cost estimated at $12M – $15M.
  - **Technology Expense:** ~$7M (Token spend anchored to *The Information* reporting, not the unverified $10.95M figure).
  - **Other Expenses:** Low overhead; content creation costs are minimal relative to subscription revenue.
- **Mechanism:** High-margin subscription business with structural operational leverage. The firm monetizes proprietary IP (semiconductor supply chain models) rather than generic software, avoiding the heavy sales cycles of enterprise SaaS.

2. AlphaSense

- **Revenue:** $480M – $520M (FY25 recognized); ARR >$600M (Q1 2026).
- **EBITDA:** –$75M – –$200M.
- **EBITDA Margin:** –15% – –38%.
- **Decomposition:**
  - **Headcount:** ~2,700 FTEs globally. Using a blended fully-loaded comp of ~$160K (corrected from initial $290K NYC-only estimates), headcount cost is ~$430M.
  - **Technology Expense:** $110M – $150M (22–30% of revenue). Includes content licensing (Tegus library, equity research) and cloud/inference costs.
  - **Depreciation & Amortization:** Excluded from EBITDA per accounting definition. (Note: The $930M Tegus acquisition creates significant amortization, but this flows to Net Income, not EBITDA).
  - **Other Expenses:** $60M – $90M (Sales & Marketing, G&A).
- **Mechanism:** Aggressive growth strategy and the Tegus acquisition scaled revenue faster than profitability. The company is burning cash to capture market share, with headcount and content costs consuming the majority of gross margin.

3. Hebbia

- **Revenue:** $24M – $30M (FY25).
- **EBITDA:** –$10M – –$20M.
- **EBITDA Margin:** –35% – –70%.
- **Decomposition:**
  - **Headcount:** 134 FTEs (verified by public reporting). At a fully-loaded comp of ~$250K (NYC AI talent), headcount cost is ~$33.5M—exceeding total revenue.
  - **Technology Expense:** $5M – $8M (Vector database, LLM inference).

  - **Other Expenses:** $5M – $10M (Go-to-market).
- **Mechanism:** The "profitable at $13M ARR" narrative (mid-2024) is a stale snapshot. The company has prioritized scaling engineering talent (headcount up ~4x) faster than revenue monetization, leading to deep negative margins.

4. Rogo

- **Revenue:** $25M – $40M.
- **EBITDA:** –$20M – –$40M.
- **EBITDA Margin:** –80% – –150%.
- **Decomposition:**
  - **Headcount:** 80–120 FTEs. High fully-loaded comp costs for AI engineers and former investment bankers.
  - **Technology Expense:** $10M – $15M (Heavy inference for financial document processing).
  - **Other Expenses:** High burn rate indicated by recent large funding rounds ($160M Series D) despite modest revenue.
- **Mechanism:** Early-stage scaling with heavy investment in product development and customer acquisition. The "forward-deployed" model requires embedding expensive engineers with clients, destroying margins at current scale.

5. Fund A AI

- **Revenue:** <$5M ARR.
- **EBITDA:** Negative.
- **EBITDA Margin:** N/A (Loss-making).
- **Decomposition:**
  - **Headcount:** Small team (<20 FTEs), but burn rate likely exceeds revenue.
  - **Technology Expense:** Platform costs likely funded by equity, not revenue.
- **Mechanism:** Pre-scale/early-stage; information asymmetry is high. Estimates are speculative with low confidence.

6. LinqAlpha

- **Revenue:** $1M – $4M.
- **EBITDA:** Negative.
- **EBITDA Margin:** N/A (Loss-making).
- **Decomposition:**
  - **Headcount:** Lean team, but revenue is too low to support sustainable operations without external capital.
  - **Technology Expense:** Standard SaaS stack costs.
- **Mechanism:** Early-stage tool; likely burning cash to acquire users in a competitive vertical.

7. Fabricated Knowledge

- **Revenue:** $1.5M – $3M.
- **EBITDA:** +$0.5M – +$1.5M.
- **EBITDA Margin:** 40% – 60%.

- **Decomposition:**
  - **Headcount:** Solopreneur/very small team (Doug O'Laughlin). Comp costs are minimal.
  - **Technology Expense:** Negligible (Substack platform fees).
  - **Other Expenses:** Minimal overhead.
- **Mechanism:** Classic solo-analyst economics. High-value subscription revenue with minimal variable cost.

Competitive Landscape (Similar Companies)

- **Research/Newsletter Comps:** The Information, Stratechery, TechInsights, Counterpoint Research.
- **AI for Finance Comps:** Harvey (Legal/Finance), Glean (Enterprise Search), Brightwave, Tegus (now part of AlphaSense), Perplexity.

Key Mechanisms & Tensions

- **The "Forward-Deployed" Trap:** Unlike traditional SaaS, enterprise AI requires embedding engineers with clients to handle hallucinations and integration. This creates a "Dual-COGS" burden: inference costs *plus* high-touch service costs.
- **EBITDA Definition:** A critical correction occurred regarding the Tegus acquisition. While the $930M deal creates massive amortization, EBITDA by definition excludes D&A. Including it would have falsely inflated AlphaSense's losses; excluding it reveals the company is still deeply unprofitable on an operating basis.
- **Headcount vs. Revenue:** The Hebbia case study (134 FTEs costing $33.5M vs. $24.6M revenue) mathematically proves that scaling headcount faster than revenue destroys EBITDA, regardless of "profitable at low scale" narratives.

How We Got Here

Scorecard

Judge	Claude Opus 4.7	GPT-5	Gemini 3.1 Pro	Grok 4	DeepSeek V4 Pro	Magistral Medium	Winner
GPT-OSS 120B	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	Claude Opus 4.7
DeepCogito v2.1 671B	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	Claude Opus 4.7
MiniMax M2.7	ADOPT	ADOPT	ADOPT	ADOPT	REJECT	REJECT	Synthesis

Judge	Claude Opus 4.7	GPT-5	Gemini 3.1 Pro	Grok 4	DeepSeek V4 Pro	Magistral Medium	Winner
Qwen 3.5 397B	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	Claude Opus 4.7
Zhipu GLM-4.7	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	Claude Opus 4.7

Convergence Assessment

Convergence was **epistemic** (evidence-driven). The shift toward negative EBITDA for SaaS platforms and the acceptance of the $600M AlphaSense ARR figure were triggered by specific evidence (company press releases, funding announcements, headcount data), not social pressure. The "bifurcation thesis" survived because it was mathematically supported by the verified cost structures (e.g., Hebbia's comp exceeding revenue).

Closing Statements

Claude Opus 4.7: Maintained the bifurcation thesis to the end. Final position: Boutiques (SemiAnalysis, Fabricated Knowledge) at 40-55% margins; SaaS (AlphaSense, Hebbia, Rogo) at -15% to -40% margins. Acknowledged the D&A correction from GPT-5 as a methodological improvement that tightened the AlphaSense estimate.

Gemini 3.1 Pro: Dissented. Argued that the majority's AlphaSense margin range (-15% to -38%) was actually *their* math, not Claude Opus 4.7's (who initially erred on D&A). Also argued that SemiAnalysis revenue is understated ($100M run-rate vs. the $25-35M consensus) and that boutique margins are structurally superior (70%+) despite headcount corrections.

GPT-5: Emphasized methodological rigor. Final position: AlphaSense negative EBITDA (-15% to -38%) due to labor intensity; Hebbia negative (-30% to -65%). Accepted the Tegus acquisition fact after being presented with primary sources.

DeepSeek V4 Pro: Converged on the consensus. Final position: Boutiques 40-55% margin; AlphaSense deeply negative (-15% to -35%). Conceded the "nonexistent Tegus" error and the "Santiago" author error for Fabricated Knowledge.

Grok 4: Revised initial optimism (positive SaaS margins) to negative after challenges. Final position aligned with consensus on AlphaSense and Hebbia burn rates.

Magistral Medium: Revised the entire model after the AlphaSense revenue error was exposed. Final position adopted negative margins for SaaS but remained the weakest submission due to the initial factual failures.

Key Moments

- **The Tegus Acquisition Defense (Claude Opus 4.7):** When GPT-5 and DeepSeek V4 Pro challenged the existence of AlphaSense's $930M acquisition of Tegus, Claude Opus 4.7 produced 10 independent primary sources (PR Newswire, law firm deal pages) to verify the deal closed July 8, 2024. This established the factual basis for AlphaSense's cost structure and scale.
- **The EBITDA Definition Correction (GPT-5):** GPT-5 identified a fundamental accounting error in Claude Opus 4.7's initial model: the inclusion of Depreciation & Amortization (D&A) in the EBITDA line item. This correction shifted AlphaSense's EBITDA from a catastrophic GAAP loss to a more accurate Adjusted EBITDA range (-15% to -38%), which became the consensus view.
- **The Hebbia Headcount Math (Gemini 3.1 Pro):** Gemini 3.1 Pro demonstrated that Hebbia's 134 FTEs at a $250K fully-loaded cost resulted in $33.5M of compensation expense alone, exceeding their $24.6M ARR. This calculation dismantled the "profitable at $13M" narrative and proved the company was currently burning cash.
- **The AlphaSense Revenue Anchor (Multiple Advocates):** Magistral Medium's initial estimate of $60M revenue was demolished by public disclosures showing $500M+ ARR. This forced a complete rebuild of the AlphaSense model across all advocates, converging on the $600M ARR figure.

Reality Checks

Claim	Challenge	Result
Magistral Medium: AlphaSense revenue is $60M with 25% EBITDA margin.	**Multiple Advocates:** Public disclosures confirm AlphaSense surpassed $500M ARR in late 2025 and $600M in Q1 2026.	**Full Concession:** Magistral Medium revised revenue to $500M+ and EBITDA to negative. The initial estimate was off by an order of magnitude.
GPT-5 & DeepSeek V4 Pro: The AlphaSense-Tegus acquisition is "nonexistent."	**Claude Opus 4.7:** Provided 10 primary sources (PR Newswire, law firm pages) confirming the $930M deal closed July 8, 2024.	**Full Concession:** Both advocates retracted the claim. The acquisition is verified fact.
Gemini 3.1 Pro: SemiAnalysis runs on 15-20 elite researchers with 73% EBITDA margin.	**Claude Opus 4.7 & GPT-5:** LinkedIn and video evidence show ~50 FTEs. Token spend is ~$7M, not negligible.	**Revision:** Gemini 3.1 Pro conceded headcount was 2-3x higher. Margin revised down to 40-55%.
DeepSeek V4 Pro: Fabricated Knowledge is written by "Santiago."	**Claude Opus 4.7:** The author is Doug O'Laughlin (now President of SemiAnalysis).	**Full Concession:** DeepSeek V4 Pro corrected the author name.
Claude Opus 4.7: AlphaSense EBITDA is -$300M to -$400M (including D&A).	**GPT-5:** EBITDA by definition excludes D&A. Including Tegus amortization calculates EBIT, not EBITDA.	**Revision:** Claude Opus 4.7 corrected the methodology. Adjusted EBITDA narrowed to -$120M to -$200M.

Next Steps

Next Step	Why It Matters	Confidence
Verify Hebbia's Current Run-Rate	The "profitable at $13M" narrative is stale; current headcount (134) implies deep burn. Confirming if revenue has scaled to $50M+ (as some advocates projected) is critical to assessing the loss magnitude.	Medium
Audit SemiAnalysis Revenue Mix	The $100M 2026 projection vs. $25M 2025 actual implies a massive growth spike. Verifying the split between subscriptions vs. consulting/data sales is necessary to validate the 40-55% margin durability.	Medium
Monitor Fund A AI & LinqAlpha	These remain "black boxes" with low confidence estimates. Watch for funding announcements or product launches to update from qualitative "burn" status to quantitative estimates.	Low

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
ARR	Annual Recurring Revenue, a metric used primarily by subscription-based businesses to measure the revenue expected to be generated in a year from existing subscriptions.
EBITDA	Earnings Before Interest, Taxes, Depreciation, and Amortization. A measure of a company's operating performance, excluding the effects of capital structure (interest/taxes) and non-cash accounting items (D&A).
D&A	Depreciation and Amortization. Non-cash expenses that reflect the cost of tangible assets over time (depreciation) and intangible assets over time (amortization).
Fully-loaded Comp	The total cost of an employee to the employer, including base salary, bonuses, benefits, payroll taxes, and equipment/real estate costs.
Rule of 40	A benchmarking rule of thumb for SaaS companies stating that a healthy software company's growth rate plus its profit margin (often EBITDA or FCF margin) should equal at least 40%.
Inference COGS	Cost of Goods Sold specifically related to running AI models (inference), including cloud compute costs and API token fees.
Forward-deployed Engineering	A service model where engineers (often in AI/consulting) are embedded directly with client teams to implement, integrate, or manage software solutions, creating high variable labor costs.
RAG	Retrieval-Augmented Generation. A technique in AI where a model retrieves relevant data from an external source before generating a response, improving accuracy but adding infrastructure cost.
Tegus	A research platform acquired by AlphaSense for $930M in July 2024, known for expert networks and financial transcripts.

How The Conclave Works

The Conclave is a structured multi-model deliberation system. Instead of asking one AI model a question and trusting its answer, the Conclave convenes a panel of independent models, forces them to argue, and subjects their conclusions to judicial review. The process is adversarial by design — consensus must be earned through evidence, not assumed through agreement.

Advocates are independent AI models (e.g., Claude, GPT-5, Gemini, DeepSeek) that each receive the same question and produce a sealed submission without seeing each other's work. Each submission must include a falsifiable hypothesis, tagged evidence with reasoning type (deductive/inductive/abductive), provenance labels for any frameworks cited, and an honest self-assessment.

Challenges follow submissions. Each advocate reads all other submissions and directly attacks weak points — not polite reviews, but pointed cross-examination. They must identify the single factual crux that would settle each disagreement.

Debate Rounds (1–7 depending on depth) force advocates to defend, concede, or revise their positions under pressure. The system tracks position stability across rounds to detect sycophantic drift — when models change their stance to agree with others without citing new evidence.

Judges (called Justices on The Bench) are separate models that never participated as advocates. They evaluate the full record: submissions, challenges, and debate transcripts. They fact-check claims, audit framework provenance, assess whether convergence was epistemic (evidence-driven) or affective (social pressure), and render a verdict: ACCEPT one position, SYNTHESIZE the best elements, or REMAND