Tribunal Session Summary

Session: conclave-20260604-011630-provide-a-best-in-f77a

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $5.2395 | Time: 34m 38s

THE PROMPT: Provide a best in class estimate of the EBITDA and EBITDA margins for Fabricated Knowledge, decomposed by top line revenue, operating expenses with technology expense, headcount, depreciation, and other expenses broken out.

THE RULING: No dollar-denominated estimate of EBITDA is possible due to the complete absence of public revenue data; however, a structural margin envelope of **67–82% Adjusted EBITDA** (pre-founder compensation) is the best-in-class estimate, derived from verifiable platform mechanics and bounded operating expense assumptions.

The Prompt

Provide a best in class estimate of the EBITDA and EBITDA margins for Fabricated Knowledge, decomposed by top line revenue, operating expenses with technology expense, headcount, depreciation, and other expenses broken out.

The Ruling

No dollar-denominated estimate of EBITDA is possible due to the complete absence of public revenue data; however, a structural margin envelope of **67–82% Adjusted EBITDA** (pre-founder compensation) is the best-in-class estimate, derived from verifiable platform mechanics and bounded operating expense assumptions.

The verdict accepts GPT-5's position (as synthesized with the converged structural view of Gemini 3.1 Pro and Claude Opus 4.7, and E) because it is the only approach that respects the evidentiary vacuum regarding revenue while still delivering actionable analytical value. The other advocates either fabricated specific revenue numbers (DeepSeek V4 Pro and Gemini 3.1 Prodvocate-B initially), refused to provide any quantitative output (Magistral Medium and Gemini 3.1 Prodvocate-D), or relied on unverified assumptions (Gemini 3.1 Pro initially). The winning position isolates the one variable that is mathematically fixed—platform fees—to construct a percentage-based margin envelope that holds true regardless of the unknown top-line revenue.

Core Analysis

The Impossibility of Dollar Estimates

The verified current facts (2024–2026) explicitly state that no public filings, investor decks, or third-party reports disclose Fabricated Knowledge’s revenue, EBITDA, margins, or cost structure. This creates a hard constraint: any specific dollar figure for revenue or EBITDA is definitionally speculative. Advocates who attempted to provide point estimates (e.g., DeepSeek V4 Pro’s $288,000 revenue scenario or Advocate-B’s $440K–$1.5M range) were forced to concede that these numbers were fabricated from unverified subscriber counts. Therefore, a "best-in-class" estimate cannot be expressed in nominal dollars.

The Structural Solution: Platform Mechanics

While the top-line revenue is unknown, the cost structure of a Substack-based business is dominated by deterministic platform mechanics that scale as a percentage of revenue. This allows for a rigorous structural decomposition of margins without knowing the absolute dollar volume.

Technology Expense (~13% of Revenue) This is the only line item that can be calculated with high confidence.

- **Substack Platform Fee:** 10% of gross subscription revenue. This is a non-negotiable standard term for the platform.
- **Payment Processing:** ~3% of gross revenue. This derives from Stripe's standard US card processing fee of 2.9% + $0.30 per transaction. On a typical annual subscription (~$220), the fixed fee is negligible (~0.14%), blending to approximately 3%.
- **Total Tech Drag:** 10% + 3% = **13%**. This is a fixed scalar cost; it applies whether the business makes $50,000 or $5,000,000.

Operating Expense Decomposition

With technology costs bounded, the remaining variability comes from operating expenses typical for a solo research business.

Other Operating Expenses (5–20% of Revenue) This category includes research data subscriptions (e.g., Bloomberg, TechInsights), software tools (analytics, email), professional services (accounting, legal), and travel (conferences like SEMICON).

- *Low End (5%):* A lean, virtual operator using free/cheap public data (SEC filings, transcripts) and minimal travel.
- *High End (20%):* A research-intensive operation with premium data terminals and heavy conference attendance.
- *Verdict:* A 5–20% band is a defensible, transparent assumption for a B2B semiconductor analyst, acknowledging the uncertainty without inventing a specific number.

Headcount (Founder Replacement Cost) The verified facts confirm no disclosed headcount beyond founder Doug O'Laughlin.

- **Cost:** A market-rate replacement salary for a senior semiconductor analyst is estimated at **$200,000–$400,000**.
- **Treatment:** This is the pivotal definitional choice.
  - **Adjusted EBITDA (SDE):** Founder compensation is added back. This reflects the cash flow available to an owner-operator.
  - **Reported EBITDA:** Founder compensation is deducted as an expense. This reflects the profitability of the entity if a third party had to hire someone to do the work.

Depreciation & Amortization (~0%) The business is asset-light (laptops, microphones). Hardware costs are typically expensed immediately or fall below capitalization thresholds. EBITDA by definition excludes D&A, so this line is effectively zero for margin analysis.

The EBITDA Margin Envelope

Combining the structural components yields the best-in-class estimate.

1. Adjusted EBITDA Margin (Pre-Founder Comp)

- Formula: 100% (Revenue) - 13% (Tech) - 5–20% (Other Opex)
- **Result: 67% – 82%**
- *Confidence:* High. This range is mathematically robust because it relies on the fixed 13% platform fee and a bounded opex assumption. It answers the "margin" question definitively without needing the revenue number.

2. Reported EBITDA Margin (Post-Founder Comp)

- Formula: Adjusted EBITDA Margin - (Founder Replacement Cost / Revenue)
- **Result: Indeterminate (Could be Negative to High)**
- *Reasoning:* Because the founder's salary is a fixed dollar cost ($200k–$400k) and the revenue is unknown, the margin impact cannot be calculated.
  - *Scenario A (Low Revenue):* If revenue is $300,000, deducting a $300,000 salary results in ~0% or negative margin.
  - *Scenario B (High Revenue):* If revenue is $1,500,000, deducting a $300,000 salary leaves a healthy ~50%+ margin.
- *Conclusion:* The verdict explicitly states that a Reported EBITDA margin cannot be provided because it requires the unknowable revenue base.

How We Got Here

Scorecard

Judge	GPT-5	Grok 4	Magistral Medium	DeepSeek V4 Pro	Claude Opus 4.7	Gemini 3.1 Pro	Winner
DeepCogito v2.1 671B	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	GPT-5
Qwen 3.5 397B	ADOPT	ADOPT	ADOPT	ADOPT	ADOPT	ADOPT	SYNTHESIZE
Zhipu GLM-4.7	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	GPT-5
Mistral Large 3	ACCEPT	REJECT	REJECT	REJECT	REJECT	REJECT	GPT-5
GPT-OSS 120B	REJECT	ADOPT	ADOPT	REJECT	PARTIAL	PARTIAL	REMAND

Convergence Assessment

The convergence was **epistemic**. The shift from dollar estimates to a structural margin envelope was triggered by the verified facts (no revenue data) and the mathematical reality of platform fees (fixed scalars). Advocates did not shift positions to agree with the majority; they shifted because the specific dollar claims (e.g., 1,200 subscribers) were proven factually impossible to defend.

Closing Statements

Gemini 3.1 Pro (Closing Statement): Gemini 3.1 Pro fully abandoned their initial dollar estimates, concluding that nominal dollar EBITDA is "epistemologically impossible." Their final position aligned with the consensus: a structural, percentage-based unit-economics envelope is the only rigorous answer. They emphasized the distinction between Adjusted EBITDA (67–82% margin) and Reported EBITDA (unresolvable), arguing that the high-margin nature of the platform mechanics is the only verifiable insight.

Claude Opus 4.7 (Closing Statement): Claude Opus 4.7 refined their position to a "structural margin decomposition," rejecting dollar estimates but defending the ratio-based approach. They argued that the brief is "partially answerable" because platform fees are verifiable constants. They highlighted that the "dollar vs. ratio" distinction was the core analytical breakthrough of the deliberation, allowing for a 67–82% pre-comp margin envelope without inventing revenue.

Magistral Medium (Closing Statement): Magistral Medium maintained a rigid refusal to estimate, arguing that even percentage-based models rely on assumptions about "Other Opex" that are not grounded in verifiable facts. They stated that any attempt to model margins would be speculative and thus violated the "best-in-class" standard defined by the constraints. They offered no quantitative output.

Grok 4 (Closing Statement): Grok 4 echoed Magistral Medium, asserting that the absence of a revenue anchor makes even percentage decompositions speculative. They argued that standard sell-side practices rely on at least one observable anchor (e.g., disclosed subscriber metrics), which are completely absent here. They concluded that no estimate of any kind is possible.

DeepSeek V4 Pro (Closing Statement): DeepSeek V4 Pro conceded their initial errors (fabricated revenue, EBITDA definition) and pivoted to the structural envelope. They argued that their final deliverable—a bounded 65–82% margin envelope—was the "best-in-class" response because it extracted maximum information from platform constants without pretending to know the unknowable top-line. They emphasized that the margin is robust to any reasonable revenue scenario.

GPT-5 (Closing Statement): GPT-5 defended their "parametric model" as the most honest deliverable. They argued that providing a bounded margin envelope (66–82%) and a pro-forma template for private inputs respects the data limitations while giving decision-makers actionable rigor. They distinguished between "knowable margin shape" (derived from fees) and "unknowable dollar magnitude" (derived from revenue), positioning their submission as the synthesis of these two truths.

Reality Checks

Claim	Challenge	Result
DeepSeek V4 Pro: "1,200 paid subscribers... yielding $288,000 top-line revenue."	**Grok 4:** "How do you defend this given the verified current facts that no public filings... disclose revenue...?"	**Full Concession.** DeepSeek V4 Pro admitted the numbers were "invented inputs" and retracted the specific dollar scenario entirely.

Claim	Challenge	Result
Claude Opus 4.7: "Fabricated Knowledge has been observed in the 'thousands' cohort... implying ~2,000–6,000 paid."	**Grok 4:** "This appears unverifiable... no public Substack leaderboard data... or pricing confirmation."	**Revision.** Claude Opus 4.7 conceded the inference was "[NOVEL]" and unverifiable, withdrawing the specific revenue range.
Gemini 3.1 Pro: "Headcount Expense: 0.0%... owner's draw / equity distribution rather than W-2 operating expense."	**Claude Opus 4.7:** "You have defined the founder's draw as 'owner's draw'... to make headcount disappear. That is an accounting choice, not a structural fact."	**Concession.** Gemini 3.1 Pro revised the model to deduct a market-rate replacement cost ($200k–$400k), acknowledging that 0% headcount inflates EBITDA artificially.
DeepSeek V4 Pro: "EBITDA (after all opex): 57,360" (after deducting $1,200 depreciation).	**Gemini 3.1 Pro:** "EBITDA stands for Earnings **BEFORE**... Depreciation... You have fundamentally violated standard accounting by deducting Depreciation."	**Correction.** DeepSeek V4 Pro conceded the definitional error and corrected the model to exclude D&A from the EBITDA calculation.

Next Steps

Next Step	Why It Matters	Confidence
Obtain private subscriber count and pricing mix	The revenue base is the only missing variable preventing a dollar-denominated EBITDA calculation.	High (Verified by facts)
Confirm specific "Other Opex" costs (data terminals, travel)	The 5–20% band is a wide assumption; actual spend determines if the margin is closer to 67% or 82%.	Medium (Industry benchmarks suggest 10-15% is likely)
Determine the valuation purpose (SDE vs. Reported EBITDA)	The "best-in-class" margin changes drastically depending on whether founder comp is treated as profit or expense.	High (Standard valuation practice)

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
Adjusted EBITDA	A metric used for owner-operator businesses where the owner's compensation is added back to the EBITDA figure to show the true cash-generating potential of the business (also known as Seller's Discretionary Earnings or SDE).
EBITDA	Earnings Before Interest, Taxes, Depreciation, and Amortization. A measure of a company's operating performance.
EBITDA Margin	EBITDA divided by Total Revenue, expressed as a percentage.
Gross Billings	The total amount of money charged to customers before any fees or refunds are deducted.
Net Revenue	Gross Billings minus platform fees (Substack), payment processing fees (Stripe), and refunds.
Platform Fee	The percentage of revenue charged by a platform (like Substack) for the use of its infrastructure and distribution network.
Reported EBITDA	The standard EBITDA metric that includes all operating expenses, including a market-rate salary for the owner/founder.
Solo-Operator	A business run by a single individual with no disclosed employees or contractors.
Substack	A paid newsletter platform that hosts content and handles billing/distribution for creators.
Stripe	A financial services platform used for processing online payments (credit card transactions).

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 for further debate.

Depth levels control rigor: T1/Spot Check (submissions only), T2/Standard Review (1 debate round, 1 judge), T3/Deep Review (3 rounds, 3 judges), T4/Full Panel (5 rounds, full panel), T5/Stress Test (5 rounds + Fresh Eyes review), T6/Red Team (7 rounds + mid-debate judicial checkpoint).

The system is deterministic code — it cannot be sycophantic. It dispatches prompts, collects responses, and enforces the adversarial structure. The models argue; the code referees.