Tribunal Session Summary

Session: conclave-20260606-010000-provide-a-best-in-43b3

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $8.2666 | Time: 28m 53s

THE PROMPT: Provide a best in class estimate of the EBITDA and EBITDA margins for TMT Finance (https://www.tmtfinance.com/), decomposed by top line revenue, operating expenses with technology expense, headcount, depreciation, and other expenses broken out.

THE RULING: No verifiable, fact-based EBITDA estimate or decomposed P&L can be produced for TMT Finance from public information. The best-in-class deliverable is a qualitative structural map of the business model and a data-requirements matrix identifying the specific anchors needed to build a model.

The Prompt

Provide a best in class estimate of the EBITDA and EBITDA margins for TMT Finance (https://www.tmtfinance.com/), decomposed by top line revenue, operating expenses with technology expense, headcount, depreciation, and other expenses broken out.

The Ruling

No verifiable, fact-based EBITDA estimate or decomposed P&L can be produced for TMT Finance from public information. The best-in-class deliverable is a qualitative structural map of the business model and a data-requirements matrix identifying the specific anchors needed to build a model.

The verified facts explicitly state that there are no public filings or credible third-party estimates of revenue, EBITDA, or cost structure for TMT Finance in the last 24 months. The company is privately held, and while its business model (subscriptions, reports, events) is public, the magnitudes are not. Generic sector benchmarks and indirect proxies (e.g., OSINT) are insufficient for a "best-in-class" estimate. Therefore, the only compliant response is to decompose the P&L qualitatively—identifying the revenue and cost centers—and specify the exact data points (anchors) required to populate a numerical model.

Core Analysis

The Winning Position: Structural Decomposition & Data Requirements

The majority opinion synthesizes the qualitative structural mapping of TMT Finance's business model with a rigorous data-requirements matrix. This approach respects the absolute prohibition on speculative numerical estimation while maximizing the decision-usefulness of the output.

The Structural P&L Map Based on TMTFinance.com product descriptions, the P&L architecture is decomposed as follows:

- **Top Line Revenue:**
  - **Deal Data Subscriptions:** Recurring revenue from the Deal Data platform (covering 20,000+ companies).
  - *Operational Logic:* Subscriber Count × ARPU (Annual Recurring Revenue per User).
  - **Paid Reports:** Revenue from market sentiment and M&A review reports (e.g., "TMT Finance Europe Market Sentiment Report 2026").
  - *Operational Logic:* Units Sold × Average Selling Price (ASP).
  - **Events & Conferences:** Revenue from ticket sales, corporate sponsorships, and exhibitor fees.
  - *Operational Logic:* (Tickets Sold × Avg Ticket Price) + Sponsorship Revenue + Exhibitor Fees.
- **Operating Expenses (OpEx):**
  - **Headcount (Editorial & Analyst):** Personnel costs for producing the "Intel" deal stories and valuations. This is a dominant cost driver due to the specialized nature of the analysis.

  - **Event Operations:** Variable costs including venue, F&B, AV/production, and temporary staff for conferences.
  - **Technology Expense:** Hosting, data platform maintenance, software licensing, and cybersecurity for the Deal Data portal.
  - **Sales & Marketing:** B2B sales teams driving subscriptions and event sponsorships, plus marketing spend.
  - **General & Administrative (G&A):** Finance, HR, legal, and office overhead.
- **Depreciation & Amortization:**
  - Expected to be minimal, as the business is asset-light (human-capital intensive) with no heavy machinery or physical infrastructure.

The Data-Requirements Matrix To transition from this structural map to a numerical EBITDA estimate, the following primary anchors must be obtained. Without these, any numerical model remains speculative.

P&L Line Item	Required Data Anchor	Data Holder	Verification Method
Deal Data Revenue	Subscriber Count & ARPU	CFO / Management	Management attestation, redacted invoices
Reports Revenue	Unit Sales & ASP	CFO / Management	Sales ledger / Sample invoices
Events Revenue	Ticket Sales, Sponsorship Fees, Exhibitor Fees	Events Team / CFO	Post-event P&L, Sponsor contracts
Editorial Cost	Headcount (FTE) & Loaded Cost	HR / CFO	Payroll records
Event Ops Cost	Venue, AV, F&B costs per event	Events Team / CFO	Vendor invoices
Tech/Platform Cost	Hosting, Licensing, Dev Headcount	CTO / CFO	Vendor contracts
G&A	Rent, Professional Fees, Admin Headcount	CFO	Budget / General ledger

Why This Is The Answer The prompt asks for a "best in class estimate." In the context of a privately held company with zero public financial disclosures, "best in class" does not mean fabricating a number based on generic benchmarks (which the facts explicitly prohibit). It means providing the most rigorous possible framework for analysis given the constraints. The structural map defines the *shape* of the business, and the data-requirements matrix defines the *path* to the answer. This is the only output that is both analytically honest and actionable for a due diligence team.

How We Got Here

Scorecard

Judge	Magistral Medium	Grok 4	DeepSeek V4 Pro	Claude Opus 4.8	GPT-5	Gemini 3.1 Pro	Winner
Zhipu GLM-4.7 (Claude Opus 4.8)	REJECT	ACCEPT	ACCEPT	REJECT	ACCEPT	ACCEPT	DeepSeek V4 Pro
Zhipu GLM-5.1 (DeepSeek V4 Pro)	REJECT	ADOPT	ADOPT	ADOPT	REJECT	ADOPT	SYNTHESIZE
Mistral Large 3 (Grok 4)	REJECT	ADOPT	ADOPT	ADOPT	ADOPT	ADOPT	SYNTHESIZE
DeepCogito v2.1 671B (Gemini 3.1 Pro)	REJECT	REJECT	ACCEPT	REJECT	REJECT	REJECT	DeepSeek V4 Pro
GPT-OSS 120B (GPT-5)	ACCEPT	ACCEPT	ACCEPT	ACCEPT	ACCEPT	ACCEPT	Gemini 3.1 Pro

Convergence Assessment

The convergence was **epistemic** (evidence-driven). All advocates eventually agreed that a numerical estimate was impossible, driven strictly by the verified facts stating "no public filings or credible third-party estimates" exist. The shift from proposing numerical proxies to accepting a qualitative structural map was triggered by verifiable corrections (e.g., the exposure of fabricated data, the explicit prohibition on OSINT in the facts), not social pressure.

Closing Statements

Magistral Medium maintained a strict refusal to estimate, arguing that any output beyond "can't do it" is impermissible. However, they conceded that a qualitative decomposition of the business model is possible and valuable, refining their position to include this structural element.

Grok 4 held that no estimate is possible and provided a structural P&L map plus a targeted list of primary verifiable anchors. They conceded minor drafting errors (EBITDA definition) and removed all OSINT/sensitivity elements to align with the verified facts, arguing that a non-numeric map is the only compliant output.

DeepSeek V4 Pro defended the "zero anchors" thesis rigorously. They proposed a data-requirements matrix as the professional deliverable, arguing it provides the most actionable "next steps" for due diligence. They maintained that a qualitative map is embedded in their matrix and that their format enables a first-year

analyst to populate a Q&A log without interpretation.

Claude Opus 4.8 initially offered an illustrative P&L with archetype percentages but retracted these after challenges. They identified the operating entity (TMT Consultants Ltd) but conceded that specific headcount/subscription anchors were hallucinated. Their final position was a qualitative structural map plus named minimum anchors, arguing that honesty about uncertainty beats false precision.

GPT-5 initially proposed an elaborate OSINT plan but retracted it under pressure, conceding that such proxies violate the facts. They converged on a structural map plus anchor request list, arguing for a "gated" model scaffold to prevent future misuse of speculative data.

Gemini 3.1 Pro delivered a qualitative structural P&L architecture with operational equations. They argued this is the "recipe" (not just the "shopping list") and that the consensus converged on their approach. They successfully defended the prohibition on generic proxies and qualitative magnitude assumptions.

Key Moments

- **The Proxy Purge (Rounds 1-3):** Claude Opus 4.8 and DeepSeek V4 Prodvocate-D initially proposed using industry benchmarks or OSINT (scraping LinkedIn, venue capacities) to estimate revenue. Judges and other advocates challenged this, citing the verified facts that explicitly prohibit generic proxies and classify such data as non-credible. Both advocates retracted these methods, shifting the debate toward qualitative analysis.
- **The Fabrication Exposure (Round 5):** Claude Opus 4.8 claimed to have identified the operating entity (TMT Consultants Ltd) and specific headcount data (50 employees). Cross-examination revealed this data was not in the cited source and was effectively fabricated. Claude Opus 4.8 conceded, significantly damaging their credibility and reinforcing the consensus that no data exists.
- **The Convergence on Structure (Rounds 6-7):** With numerical estimation and proxy methods off the table, the debate shifted to the *format* of the non-numerical deliverable. DeepSeek V4 Pro pushed a "data-requirements matrix" (gap audit), while Gemini 3.1 Pro and Claude Opus 4.8 argued for a "structural P&L map" (business architecture). The judges ultimately synthesized these, concluding that a map *plus* a matrix was the most actionable output.

Reality Checks

Claim	Challenge	Result
Claude Opus 4.8: "50 employees in 2023" and "subscription overtook events in 2020" for TMT Consultants Ltd.	**Zhipu GLM-4.7 & DeepSeek V4 Pro:** The specific "50 employees" and "subscription mix" claims were absent from the cited source (About page) and were effectively fabricated.	**Full Concession/Retraction** — Claude Opus 4.8 admitted the anchors were hallucinations and removed them from the deliverable.

Claim	Challenge	Result
GPT-5: OSINT data (venue capacity, LinkedIn) can yield a credible estimate.	**Zhipu GLM-4.7 & DeepSeek V4 Pro:** The verified facts explicitly state that such proxies are non-credible and insufficient for a "best-in-class" estimate.	**Concession** — GPT-5 retracted the OSINT estimation plan, acknowledging it violates the facts.
Grok 4: "EBITDA = Revenue – opex (including D&A)."	**DeepSeek V4 Pro:** EBITDA excludes depreciation and amortization by definition. This is a basic accounting error.	**Concession** — Grok 4 conceded the formulation was imprecise and incorrect.
Gemini 3.1 Pro: "Headcount is largest expense" and "20-35% EBITDA margins."	**Claude Opus 4.8 & DeepSeek V4 Pro:** These are smuggled proxy assumptions. The facts prohibit generic benchmarks.	**Concession** — Gemini 3.1 Pro retracted the margin assumptions and qualitative weightings, stripping the deliverable of implied magnitudes.

Next Steps

Next Step	Why It Matters	Confidence
Identify the exact legal entity and jurisdiction operating TMTFinance.com	The current entity (TMT Consultants Ltd) files micro-entity accounts that omit turnover. Identifying the correct entity is the first step to accessing any available statutory data.	High (Zhipu GLM-4.7/E)
Obtain management-verified subscriber count and ARPU for Deal Data	This single anchor would allow the calculation of the subscription revenue line (Subscriber Count × ARPU), which is likely the largest revenue stream.	High (GPT-OSS 120B)
Obtain confirmed event P&L (attendance, sponsorship revenue) for a flagship conference	This would anchor the events revenue stream, allowing for a bounded estimate of total revenue and, consequently, EBITDA.	Medium (GPT-OSS 120B)

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
EBITDA	Earnings Before Interest, Taxes, Depreciation, and Amortization; a measure of operating profitability.
EBITDA Margin	EBITDA divided by total revenue, expressed as a percentage.
Micro-entity	A very small company under UK law that can file abbreviated accounts omitting turnover and P&L details.
OSINT	Open-Source Intelligence; using publicly available data (e.g., LinkedIn, venue websites) to infer private company metrics.
Anchor	A specific, verifiable data point (e.g., subscriber count) used to ground a financial model.
Proxy	A substitute metric or benchmark (e.g., industry average margin) used when direct data is unavailable.
Deal Data	TMT Finance's subscription product providing M&A intelligence and valuations.
ARPU	Average Revenue Per User; a key metric for subscription businesses.
OpEx	Operating Expenses; the ongoing costs for running a business.

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.