Tribunal Session Summary

Session: conclave-20260514-205616-after-its-ipo-today-0fa5

Depth: T6

Advocates: 6 | Judges: 5 | Cost: $5.7493 | Time: 23m 45s

THE PROMPT: After its IPO today with a closing market cap of $67B, will Cerebras be worth materially more, less, or about the same in two years? Please include a probability estimate, key factors influencing the valuation, and what developments could significantly alter this outlook.

THE RULING: Cerebras will likely be worth **about the same or slightly less** than its $67B IPO market cap in two years, with a **~55% probability of material decline** (>25% down), **~30% chance of flat valuation**, and **~15% upside potential**.

The Prompt

After its IPO today with a closing market cap of $67B, will Cerebras be worth materially more, less, or about the same in two years? Please include a probability estimate, key factors influencing the valuation, and what developments could significantly alter this outlook.

The Ruling

Cerebras will likely be worth **about the same or slightly less** than its $67B IPO market cap in two years, with a **~55% probability of material decline** (>25% down), **~30% chance of flat valuation**, and **~15% upside potential**.

This outcome is driven not by a failure to execute on its massive backlog, but by **valuation de-rating** due to extreme customer concentration and a visible terminal threat from OpenAI's custom silicon development. The $24.6B backlog and $20B+ OpenAI contract provide strong near-term revenue visibility, but the market will apply a severe "Single-Tenant Monopsony Discount" as OpenAI's in-house Broadcom partnership advances. At a $67B valuation, Cerebras trades at a ~36× forward sales multiple on projected 2027 revenue, which is unsustainable for a hardware vendor with >80% of its revenue tied to a single customer actively building a replacement. The primary risk is multiple compression, not revenue shortfalls.

The strongest evidence is the combination of verified facts: the 86% revenue concentration, the OpenAI-Broadcom partnership announced October 13, 2025, with a 2H2026 production target, and the historical precedent of hardware vendors facing insourcing by dominant customers (e.g., Lattice, Habana) suffering severe multiple compression. The main caveat is the reported $18B financing snag in the Broadcom project, which could delay the substitution threat and preserve the multiple.

Core Analysis

The $24.6B Backlog: A Revenue Floor, Not a Growth Engine

Cerebras enters the public market with an exceptional degree of revenue visibility. The **$24.6B backlog** reported at the end of 2025, combined with the **expanded OpenAI contract** valued at more than $20B through 2028, creates a powerful floor for near-term financial performance. This backlog represents **48× Cerebras's 2025 revenue of $510M**, a ratio that is extraordinarily high for any capital equipment or hardware company. This is not pipeline or aspiration; it is legally contracted revenue.

However, the *timing* of this conversion is critical. Early bullish arguments assumed a rapid, front-loaded conversion, projecting $6–8B in annualized revenue. This was based on an arithmetic average over a 3–4 year horizon. The verified facts, however, provide a more conservative guide. Reporting indicates that Cerebras expects to recognize **15% of its backlog cumulatively in 2026 and 2027**. This translates to approximately **$3.7B in revenue over two years**, or an average of **~$1.85B per year**. This is a significant

increase from 2025, but it is not the explosive growth required to justify a $67B valuation.

This slower, more realistic conversion curve implies that Cerebras's 2028 revenue will be in the **$3–5B range**, assuming the back-loaded years of the contract are fulfilled. This provides a clear, defensible revenue floor. The company's $4.2–4.8B in IPO proceeds (from 28M shares at $150–160) provides ample capital to fund the necessary manufacturing, inventory, and deployment, reducing the risk of a supply-chain chokepoint. The backlog is real and de-risks near-term execution, but it does not, by itself, create a path to material upside.

The Terminal Threat: OpenAI's In-House Silicon and the "Single-Tenant Monopsony Discount"

The dominant factor in Cerebras's two-year valuation is not revenue growth, but **multiple compression**. The $67B market cap implies a **~36× forward sales multiple** on the projected $1.85B 2027 revenue. This multiple is more typical of a high-growth software company with a diversified customer base and a strong moat. Cerebras is a hardware vendor with **86% of its 2025 revenue coming from just two customers**, with OpenAI as the dominant anchor.

This concentration creates a structural vulnerability. The market has a well-documented history of applying a severe discount to hardware suppliers who are functionally captive to a single, monopsonistic buyer. This is known in equity research as a **"Single-Tenant Monopsony Discount."** The discount is not applied only when the customer leaves; it is applied when the *risk* of departure becomes visible.

That risk is now crystal clear. On **October 13, 2025**, OpenAI and Broadcom announced a multi-year partnership to deploy **10 gigawatts (GW) of custom AI accelerators**, with production targeted to start in the **second half of 2026**. This is not a vague aspiration; it is a primary-source, dated announcement. Cerebras's 750 megawatts (MW) of compute for OpenAI is **less than 3%** of this 10 GW in-house program. The scale of the Broadcom project is **13.3 times larger** than Cerebras's entire OpenAI commitment.

Markets are forward-looking discounting mechanisms. They do not wait for the final invoice to be paid before re-rating a stock. The announcement of a 10 GW, in-house alternative to Cerebras's 750 MW contract is a terminal threat. Even if the project faces delays—such as the reported **$18B financing snag**—the direction is unambiguous. By May 2028, the market will have had **18–24 months** of visibility into OpenAI's progress. This is more than enough time for investors to price in a significant de-rating.

Historical precedent supports this. Hardware companies like Lattice and Habana, which were heavily dependent on a single hyperscaler, saw their multiples collapse when their customers began insourcing. The market will not pay a 36× multiple for a company whose primary revenue stream is on a clear path to obsolescence.

The AWS Channel: Optionality, Not a Diversification Solution

The planned **AWS distribution path**, with CS-3 systems to be deployed in AWS data centers and offered through **Amazon Bedrock**, is a positive development. It provides a potential channel to diversify Cerebras's customer base beyond OpenAI. However, it is not a near-term solution to the concentration problem.

Historical evidence shows that AWS's support for third-party accelerators is often a stopgap. Partners like **Habana** and **Graphcore** were listed on AWS but failed to gain meaningful market share against AWS's own **Trainium** and **Inferentia** silicon. AWS has a documented pattern of using partner deployments to inform its own roadmap before eventually displacing them. The AWS partnership is therefore best viewed as **optionality**—a potential upside driver after 2028, not a base-case revenue stream that can offset the

OpenAI concentration risk in the next two years.

Competitive and Manufacturing Risks Are Secondary

Claims of a **2–3× performance advantage** for Cerebras's wafer-scale engine (WSE) architecture over Nvidia's Blackwell and Rubin platforms are unverified. No independent, current-generation benchmarks (e.g., MLPerf) have been produced to substantiate this claim. While the architecture offers theoretical benefits in on-chip bandwidth and latency, the lack of public validation means this moat cannot be relied upon as a valuation driver.

Similarly, concerns about **TSMC's ability to scale wafer-scale production** were overstated. Early claims that a 20–40× scale-up was required were based on a misreading of the backlog conversion timeline. The actual requirement is a **~3.6× increase** from the $510M 2025 base, which is a challenging but physically plausible ramp, especially with $4.2–4.8B in new capital. This is an execution risk, not a valuation driver.

The Capital Structure and Lockup Dynamics

The IPO raised **~$4.2–4.8B** in gross proceeds, not the $5.55B claimed by some advocates. This provides a strong war chest to fund operations and mitigate working capital constraints. However, this does not change the fundamental valuation story. The capital markets are also subject to mechanical forces. The IPO was **more than 20x oversubscribed**, leading to a significant price increase from the initial range. This pattern, seen in IPOs like Arm and Klaviyo, is often followed by a **20–45% drawdown** after the 180-day lockup expires, as insider supply meets saturated demand. This mechanical headwind will add to the multiple compression pressure.

Probability Estimate and Scenario Outlook

Based on the synthesis of the strongest evidence, the probability distribution for Cerebras's valuation in two years is:

- **55% probability of materially less (<$50B, >25% decline)**: This is the most likely outcome. It requires no execution failure. Even if Cerebras delivers $3–5B in 2028 revenue, a multiple compression to **10–12× forward sales** (a historically normal range for a hardware vendor with terminal risk) would result in a $30–50B valuation.
- **30% probability of about the same ($50–67B)**: This requires a slower-than-expected ramp of the Broadcom silicon, perhaps due to the financing snag, or meaningful traction from the AWS channel and new customer wins that reduce concentration below 60%.
- **15% probability of materially more (>$87B, >30% up)**: This is a low-probability, high-upside scenario. It requires Cerebras to successfully diversify its customer base (e.g., a new >$1B customer), demonstrate a clear 2×+ performance advantage over Blackwell/Rubin in independent benchmarks, and see the AWS channel generate significant third-party revenue.

What Could Significantly Alter This Outlook

- **Upside Catalysts**:
  - **Delays in OpenAI-Broadcom silicon**: A 12+ month delay in the 2H2026 production target would push the substitution threat beyond the two-year horizon, preserving the multiple.
  - **AWS Bedrock adoption by major third parties**: If Cerebras becomes a default option for non-OpenAI workloads on AWS, it could fundamentally reduce concentration risk.

  - **New large customer wins**: A contract with another hyperscaler (e.g., Meta, Anthropic) worth >$500M would validate the technology and diversify the book.
  - **Independent performance validation**: MLPerf results showing a 2×+ advantage over Blackwell would support pricing power and margin durability.

- **Downside Catalysts**:
  - **Accelerated Broadcom ramp**: If the $18B financing is secured quickly and production begins on schedule, the multiple compression will be faster and more severe.
  - **OpenAI contract amendment**: Any public disclosure that reduces the scope or value of the OpenAI commitment would trigger an immediate sell-off.
  - **Manufacturing or yield issues**: Any delay in CS-3 shipments would confirm execution risk and accelerate the de-rating.
  - **Broader AI capex pullback**: A macroeconomic shock or regulatory change that reduces AI spending would disproportionately affect a concentrated, high-multiple name like Cerebras.

How We Got Here

Scorecard

Judge	DeepSeek V4 Pro	GPT-5	Magistral Medium	Grok 4	Gemini 3.1 Pro	Claude Opus 4.7	Winner
GPT-OSS 120B	REJECT	REJECT	REJECT	REJECT	REJECT	**ACCEPT**	Claude Opus 4.7
Qwen 3 235B (Cerebras)	**ACCEPT**	REJECT	REJECT	REJECT	REJECT	REJECT	DeepSeek V4 Pro
Qwen 3.5 397B	REJECT	REJECT	REJECT	REJECT	REJECT	**ACCEPT**	Claude Opus 4.7
DeepCogito v2.1 671B	REJECT	REJECT	REJECT	REJECT	REJECT	**ACCEPT**	Claude Opus 4.7
Zhipu GLM-5	REJECT	REJECT	REJECT	REJECT	REJECT	**SYNTHESIZE**	Claude Opus 4.7

Convergence Assessment

The convergence was **epistemic**, driven by verifiable corrections and new evidence. The most significant shifts were triggered by hard data: the **retraction of the $5.55B cash figure**, the **concession on the AWS revenue assumption**, and the **introduction of the primary-source Broadcom announcement**. These were not social pressures or sycophantic drift; they were forced by the adversarial system. For example, GPT-5 significantly revised their position from a 60% probability of material upside to a modal "about the same" after engaging with the OpenAI-Broadcom evidence. This was a clear, evidence-based shift in belief, not a move to conform.

Closing Statements

DeepSeek V4 Pro In their closing statement, DeepSeek V4 Pro maintained a bullish position, assigning a 50% probability to a materially higher valuation. Their core argument was that the $24.6B backlog and $20B+ OpenAI commitment create a "revenue floor" that compels a higher valuation, even after conceding the AWS and performance moat claims. What changed from their opening was a significant refinement: they abandoned their initial 60% probability and the $100–200B valuation target, acknowledging the extreme starting multiple and concentration risk. They found the most challenging arguments to be those from Claude Opus 4.7 and Grok 4dvocate-E on the terminal risk from OpenAI's insourcing. Their strongest surviving evidence was the backlog itself, which they argued was a "mathematical" floor that could not be ignored. They also highlighted the **financing snag in the Broadcom project** as a new piece of evidence that weakened the bear case.

GPT-5 GPT-5’s final position was a significant shift from their opening. They ended with a 25% probability of material upside, 45% for flat, and 30% for material downside, making "about the same" their modal outcome. This was a major revision from their initial 60% probability of material upside. The change was driven by the evidence of the OpenAI-Broadcom partnership, which forced them to accept that the market would forward-price a 2029–2030 run-off of the OpenAI contract. The most challenging arguments were from Gemini 3.1 Pro, who demanded a historical precedent for a hardware company maintaining a >25x forward sales multiple while its dominant customer ramped in-house silicon—a precedent that does not exist. Their strongest surviving evidence was the backlog and named commitments, which they argued provided "unusual revenue visibility," and the oversubscribed IPO, which "derisks supply-chain execution."

Magistral Medium Magistral Medium’s closing statement showed no meaningful change from their opening submission. They maintained a 60% probability of material upside, citing the same generic factors: strong AI market demand, the $24.6B backlog, and strategic partnerships. They acknowledged the challenges on customer concentration and generic growth trends but did not revise their probability or provide a quantitative bridge from revenue to valuation. Their strongest surviving evidence remained the broad AI market growth trend and the existence of the backlog. They failed to engage with the specific financial mechanics or competitive benchmarks that dominated the debate, making their position the least responsive to the adversarial process.

Grok 4 Grok 4’s position was refined but directionally consistent. They ended with a 48% probability of material upside, down from 65%. The change was driven by the retraction of their unverified claims: the $5.55B cash figure, the 15% annual backlog conversion, and the performance moat. The most challenging arguments were from multiple advocates on the arithmetic inconsistency of their probability weighting. Their strongest surviving evidence was the $24.6B backlog and the AWS distribution agreement, which they

argued "removes the go-to-market constraint." They also cited the absence of any public indication that OpenAI intends to reduce its commitment.

Gemini 3.1 Pro Gemini 3.1 Pro's final position was unchanged: a 58% probability of material downside. Their framework evolved from a focus on physical manufacturing constraints to a refined argument centered on "multiple compression and terminal substitution risk." The most challenging arguments were from Claude Opus 4.7 on the lack of a historical precedent for a hardware company maintaining a high multiple during insourcing, which they conceded was a valid point. Their strongest surviving evidence was the **inescapable math of the Broadcom/OpenAI substitution timeline**, the **86% revenue concentration**, and the **historical iron law of hardware multiple compression**.

Claude Opus 4.7 Claude Opus 4.7's final position was refined but directionally consistent: a 55% probability of material downside. They conceded precision overreaches on the 33GW substitutability claim and the 5–10x float expansion, but their core thesis of multiple compression due to visible terminal risk remained intact. The most challenging arguments were from DeepSeek V4 Pro on the back-loaded revenue model, which they countered by exposing the arithmetic inconsistency in the bull case. Their strongest surviving evidence was the **arithmetic of multiple compression**, the **customer concentration reference class**, and the **forward-pricing mechanism for terminal risk**.

Reality Checks

Claim	Challenge	Result
Grok 4: "Post-IPO balance sheet ($5.55B cash raised) removes capital constraints."	**DeepSeek V4 Pro:** "Show your math. 28M shares at $150–160 = $4.2–4.5B gross, before fees. You've inflated the war chest by ~25%."	Grok 4 fully conceded the error, retracting the $5.55B figure and acknowledging the actual proceeds were ~$4.2–4.8B. This demolished a key pillar of the bull case.
DeepSeek V4 Pro: "Even assuming a conversion horizon of 3–4 years, that implies annualized revenue of $6–8B from existing contracts alone."	**GPT-5:** "How do you defend this given the prospectus-summary reporting that management expects to recognize only '15% of [the $24.6B backlog] in the years 2026 and 2027' (interpreted as cumulative, not per year)?"	DeepSeek V4 Pro defended the arithmetic average but conceded that the 15% figure was cumulative, not annual. This forced a revision of the revenue timeline and weakened the bull case's near-term growth narrative.
Gemini 3.1 Pro: "Fulfilling a $24.6B backlog by 2028 requires a ~20x to 40x annualized scale-up in manufacturing."	**DeepSeek V4 Pro:** "You've conflated the total backlog with the near-term revenue required to avoid a collapse. The 15% cumulative guidance means only ~$3.7B needs to convert by end-2027, a 3.6x ramp, not 20–40x."	Gemini 3.1 Pro fully conceded the error, acknowledging that the physical constraint was far less binary than claimed. This was a dramatic correction that strengthened the bear case's focus on multiple compression over supply chain risk.

Claim	Challenge	Result
Claude Opus 4.7: "Median 24-month return for >50x revenue semi IPOs since 2015 is approximately -35%."	**Gemini 3.1 Pro:** "This appears unverifiable, as you explicitly admit it is an estimate you 'constructed from memory.' Using a self-admittedly unaudited, remembered statistic as the mathematical foundation for a structural valuation argument is unacceptable."	Claude Opus 4.7 conceded the point, retracting the specific -35% figure and replacing it with a qualitative argument about the historical tendency for high-multiple semi IPOs to compress. This was a significant correction of a load-bearing claim.

Next Steps

Next Step	Why It Matters	Confidence
Monitor OpenAI-Broadcom silicon progress for signs of delay (e.g., funding, tape-out)	A 6–12 month delay would defer the terminal risk, supporting a higher valuation multiple.	~50%
Track AWS Bedrock for disclosures of third-party Cerebras customers and utilization	This is the primary path to diversification; traction would reduce concentration risk.	~30%
Await independent performance benchmarks (e.g., MLPerf) comparing CS-3 to Nvidia Blackwell/Rubin	This would validate or invalidate the core competitive moat and pricing power.	~40%
Analyze quarterly filings for changes in top customer concentration and backlog conversion rate	This provides real-time data on the two most critical valuation drivers.	~70%
Assess the impact of the IPO lockup expiration in November 2026 on share price	This mechanical supply event could create a short-term headwind.	~80%

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
Backlog	The total value of signed, non-cancelable customer contracts for future delivery of goods or services, representing future revenue visibility.
Customer Concentration	A situation where a large portion of a company's revenue comes from a small number of customers, creating significant business risk.
Forward Sales Multiple	A valuation metric calculated as Enterprise Value divided by projected revenue for the next twelve months (NTM).
In-House Silicon	Custom-designed semiconductor chips developed and used internally by a company, as opposed to being purchased from a merchant supplier.
IPO Lockup	A contractual restriction that prevents insiders and early investors from selling their shares for a specified period (often 180 days) after an initial public offering.
Monopsony	A market structure where there is only one buyer for a product or service, giving that buyer significant power over suppliers.
Multiple Compression	The process by which a company's valuation multiple (e.g., P/E, EV/S) decreases, often due to reduced growth expectations or increased risk.
Oversubscribed	An IPO where the demand for shares from investors exceeds the number of shares available for sale.
Post-IPO	The period after a company has completed its initial public offering and its shares are trading on a public exchange.
Reference-Class Forecasting	A method of predicting the outcome of a specific case by comparing it to the outcomes of a group of similar past cases.
Revenue Recognition	The accounting process of recording revenue in the financial statements when it is earned and realizable.
Trailing Revenue	A company's revenue over the most recent twelve-month period.
Wafer-Scale Engine (WSE)	A type of semiconductor architecture where an entire silicon wafer is used as a single, massive chip, as opposed to dicing the wafer into many smaller dies.

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.