Tribunal Session Summary

Session: conclave-20260512-144435-how-many-cava-restaurants-e116

Depth: T6

Advocates: 2 | Judges: 2 | Cost: $1.0373 | Time: 34m 15s

THE PROMPT: How many CAVA restaurants will CAVA Group, Inc. (NYSE: CAVA) report having open at the end of Q1 fiscal 2026, in its earnings release on May 19, 2026?

THE RULING: The tribunal cannot definitively determine how many CAVA restaurants will be reported at the end of Q1 fiscal 2026. The case is **REMANDED** for primary-source verification due to insufficient evidence. While both advocates converged on **512 restaurants**—based on a secondary source claiming CAVA opened **73 net new locations in Q1 FY2026**—this figure lacks confirmation from CAVA's official May 19, 2026, Form 8-K, press release, or earnings transcript. The claim of 73 openings represents a **fivefold acceleration** over historical Q1 cadence (14–15 units in prior years), which contradicts CAVA's typical back-loaded development pattern and demands direct validation.

The Prompt

How many CAVA restaurants will CAVA Group, Inc. (NYSE: CAVA) report having open at the end of Q1 fiscal 2026, in its earnings release on May 19, 2026?

The Ruling

The tribunal cannot definitively determine how many CAVA restaurants will be reported at the end of Q1 fiscal 2026. The case is **REMANDED** for primary-source verification due to insufficient evidence.

While both advocates converged on **512 restaurants**—based on a secondary source claiming CAVA opened **73 net new locations in Q1 FY2026**—this figure lacks confirmation from CAVA's official May 19, 2026, Form 8-K, press release, or earnings transcript. The claim of 73 openings represents a **fivefold acceleration** over historical Q1 cadence (14–15 units in prior years), which contradicts CAVA's typical back-loaded development pattern and demands direct validation.

The strongest evidence—CAVA's FY2024 ending count of **367 restaurants** and FY2025 guidance of **68–70 net openings**—is verified via primary sources. However, the critical Q1 FY2026 input (73 net opens) rests solely on an unverified Octagon AI summary. Until the primary disclosure is retrieved, the tribunal cannot accept convergence on an unconfirmed number, no matter how internally consistent it appears.

Core Analysis

The Converged Estimate: 512 Restaurants

Both advocates—Claude Opus 4.7 and GPT-5—ultimately arrived at the same point estimate: **512 CAVA restaurants open at the end of Q1 fiscal 2026**. This number is derived from a coherent arithmetic chain:

- **FY2024 ending count**: 367 restaurants (reported in CAVA's 10-K, confirmed primary source)
- **FY2025 net openings**: ~72 (implied by 512 – 73 = 439 end-2025; 439 – 367 = 72 net adds)
- **Q1 FY2026 net openings**: 73 (cited in Octagon AI summary of May 19, 2026 earnings release)
- **Q1 FY2026 ending count**: **512** (439 + 73)

This chain is mathematically sound and directionally consistent with CAVA's growth trajectory. The implied FY2025 net adds (~72) slightly exceed the last-disclosed guidance of 68–70 but remain within plausible variance. The full-year FY2026 guidance of **74–76 net openings**—also cited in the Octagon AI summary—makes a Q1 contribution of 73 highly front-loaded but not impossible.

Anchoring on Verified Milestones

The analysis is grounded in several **verified primary-source anchors**:

- **309 restaurants at FY2023 end**: Disclosed in CAVA's 2023 Form 10-K. This is the starting point for any forward projection.
- **367 restaurants at FY2024 end**: Confirmed in CAVA's Q4 FY2024 earnings release (February 2025). This corrects an early error by GPT-5, who initially used outdated guidance (~50 net adds) instead of actuals (58 net adds).
- **68–70 net openings guided for FY2025**: Stated in CAVA's Q3 FY2025 earnings call (November 2025). This sets a firm bound on the FY2025 ending base: 367 + 68 = **435**, 367 + 70 = **437**.
- **Q1 is a 16-week quarter**: Confirmed in CAVA's fiscal calendar disclosures. This structural detail invalidates peer-based heuristics (e.g., "Q1 = 20–25% of annual openings") unless normalized for duration.

These anchors are **uncontested and verified** across both judicial opinions and the Fresh Eyes review.

The Anomalous Q1 FY2026 Surge

The central tension in the analysis is the claim that CAVA opened **73 net new restaurants in Q1 FY2026**—a figure that **represents a 5x acceleration** over historical Q1 cadence:

- **Q1 FY2024**: 14 net openings (309 → 323)
- **Q1 FY2025**: 15 net openings (367 → 382)
- **Q1 FY2026**: **73 net openings** (439 → 512)

This surge is operationally extreme. Fast-casual chains like CAVA typically back-load openings due to permitting, construction timelines, and lease finalizations. A reversal of this pattern would constitute a **material shift in operating rhythm** and would almost certainly be highlighted in management commentary.

The **Octagon AI summary**—the sole source for the 73 figure—provides no context for this anomaly. Was it driven by:

- A **strategic front-loading** to lock in Midwest expansion (e.g., Detroit, Indianapolis)?
- A **pipeline backlog** clearing in Q1?
- **Inorganic additions** (conversions, acquisitions)?
- Or a **data error**?

Without the primary 8-K or earnings transcript, these questions remain unanswerable.

Source Hierarchy and Risk Assessment

The tribunal applied a strict source hierarchy:

Source Type	Examples	Reliability
Primary	10-K, 8-K, earnings call transcript	High — audited, SEC-filed
Secondary	Press releases, investor presentations	Medium — official but unaudited
Tertiary	Aggregator summaries (e.g., Octagon AI)	Low — unverified, risk of misreading

The **512 estimate** is built on:

- **Primary**: FY2024 count (367), FY2025 guidance (68–70), fiscal calendar
- **Secondary**: Historical Q1 openings (14, 15)
- **Tertiary**: Q1 FY2026 openings (73), FY2026 guidance (74–76)

The **critical input (73)** is tertiary. While the arithmetic chain is internally consistent, **consistency does not equal truth**. The tribunal cannot accept a financial forecast based on a single unverified aggregator, especially when it contradicts historical patterns.

Confidence Calibration

Both advocates acknowledged uncertainty:

- **Claude Opus 4.7**: ~80% confidence on 512, ±3 unit band due to aggregator risk
- **GPT-5**: ~80–85% confidence, same band

This calibration is appropriate. The **509–515 range** is robust: even if the aggregator rounded or conflated company-operated vs. total units, the magnitude (~510s) is supported by the implied FY2025 base (~439) and the front-loaded guidance (74–76).

However, **80% confidence is not sufficient for a definitive ruling** in a financial forecasting context. The tribunal requires **near-certainty** for acceptance, especially when anomalies are present.

How We Got Here

Scorecard

Judge	Claude Opus 4.7	GPT-5	Winner
Qwen 3 235B	ACCEPT	REJECT	Claude Opus 4.7
Qwen 3 235B	REMAND	REMAND	REMAND

Convergence Assessment

Convergence was **epistemic**, not affective. Both advocates shifted positions in response to **new evidence**—the Octagon AI summary—not social pressure. Claude Opus 4.7 resisted initially due to the anomaly but revised after verifying internal consistency. GPT-5 revised more quickly but acknowledged source risk. Position stability data shows both models abandoned their original estimates only after evidence review, not peer influence.

Closing Statements

Claude Opus 4.7: My final position is **512 restaurants**, with a 509–515 range. I arrived here not through my original trajectory model (which projected ~450) but because a third-party aggregator appeared to summarize the actual earnings release. While the 73-opening surge is anomalous, the arithmetic chain (367 → 439 → 512) is internally consistent. I hold ~80% confidence due to aggregator-source risk. The correct framework is evidentiary: go to the highest-quality available source, weight it by reliability, and flag uncertainty. Only primary-source inspection can resolve the residual doubt.

GPT-5: I close with **512 restaurants** as the most defensible answer. My initial 433 was built on outdated guidance and peer heuristics. The Octagon AI summary provided a specific, testable claim that reconciled with known anchors: 367 (FY2024) → 439 (FY2025) → 512 (Q1 FY2026). The 16-week Q1 makes the 73-opening surge feasible, and the 1,000-by-2032 target supports continued acceleration. I ascribe ~80–85% confidence, pending primary-source pull. The absence of a competing, credible alternative strengthens the case for 512, but the ±3 unit band must be acknowledged.

Reality Checks

Claim	Challenge	Result
GPT-5: CAVA ended FY2024 with ~359 restaurants	**Claude Opus 4.7:** That's factually wrong — CAVA reported 367 in its 10-K, a 58-unit net add	Full concession: GPT-5 revised to 367 and acknowledged the error
GPT-5: CAVA would open ~60 units in FY2025	**Claude Opus 4.7:** Management guided to 68–70, not 60 — you substituted your own "conservative" number	Full concession: GPT-5 accepted the 68–70 guidance and revised its FY2025 base
Claude Opus 4.7: CAVA would open ~15 units in Q1 FY2026	**GPT-5:** The Octagon AI summary claims 73 net openings — a fivefold acceleration	Full revision: Claude Opus 4.7 abandoned 450 and adopted 512 after verifying internal consistency

Next Steps

Next Step	Why It Matters	Confidence
Retrieve CAVA's Form 8-K or press release from May 19, 2026	This is the only way to verify the 512 restaurant count and 73 net openings claim	100%

Next Step	Why It Matters	Confidence
Confirm the exact number of restaurants open at Q1 FY2026 end	Resolves the core question with primary evidence	100%
Cross-check whether Q1 FY2026 openings include conversions or acquisitions	Ensures the 1,000-by-2032 organic growth target remains valid	90%
Review the earnings call transcript for commentary on front-loading	Explains the operational anomaly of 73 Q1 openings	85%
Validate the Octagon AI summary's accuracy and track record	Assesses the reliability of third-party aggregators in financial forecasting	70%

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
10-K	Annual financial report filed with the SEC, containing audited financials and operational disclosures.
8-K	Current report filed with the SEC to announce material events, such as earnings releases.
Fiscal Q1	The first quarter of a company’s fiscal year; for CAVA, a 16-week period.
Net new restaurant openings	The number of new locations opened minus any closures in a given period.
Primary source	An original, authoritative document such as an SEC filing or earnings transcript.
Secondary source	A report or summary derived from primary sources, such as a press release.
Tertiary source	A third-party aggregation or interpretation of secondary sources, such as an analyst summary.
Shared-data hallucination	A risk in synthetic forecasting where multiple models converge on an incorrect number due to reliance on the same unverified source.

How The Conclave Works

The Tribunal is a structured multi-model deliberation system. Instead of asking one AI model a question and trusting its answer, the Tribunal 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.