Tribunal Session Summary

Session: conclave-20260507-140313-can-a-pe-style-89b9

Depth: T6

Advocates: 5 | Judges: 4 | Cost: $3.6987 | Time: 21m 18s

THE PROMPT: Can a PE-style roll-up of fragmented outpatient radiology centers clear a 20% IRR over a 5-year hold? Yes or no, with conviction level. To answer it, work through: Structural choice — which ownership model gives the best shot at the return: (a) own imaging sites and contract with independent radiology groups, (b) own the radiology groups, or (c) vertically integrate both? Evaluate each on margin capture, operational complexity, regulatory exposure (Stark, anti-kickback, corporate practice of medicine), reimbursement split between technical and professional components, and exit optionality at year 5. AI disruption overlay — over the same 5-year window, how does AI in radiology (auto-detection, triage, workflow automation, read-volume compression) affect the return math? Does it commoditize the read and push value to the site/technical side, or does owning the read workflow become the durable moat? Factor in second-order effects on radiologist labor cost, throughput, and payer behavior. Return math — under base / bull / bear scenarios, what IRR does the recommended structure actually produce? What entry multiple, organic growth rate, margin expansion, and exit multiple are required to clear 20%, and how realistic are those assumptions given current market conditions? Output I want: 1 Yes or no — does this clear 20% over 5 years, and at what confidence level? 2 The recommended structure and the IRR it produces in the base case 3 The top three risks that would push the return below 20% 4 The one variable that matters most to whether the thesis works

THE RULING: Partially — a PE-style roll-up of fragmented outpatient radiology centers can clear a 20% IRR over a 5-year hold, but only under a site-led, technically focused model with embedded contractual alignment that captures value from throughput and operational leverage, while avoiding the regulatory and execution risks of full vertical integration. The base case IRR is 18–22%, contingent on realistic assumptions for margin expansion and modest multiple expansion.

The Prompt

Can a PE-style roll-up of fragmented outpatient radiology centers clear a 20% IRR over a 5-year hold? Yes or no, with conviction level. To answer it, work through: Structural choice — which ownership model gives the best shot at the return: (a) own imaging sites and contract with independent radiology groups, (b) own the radiology groups, or (c) vertically integrate both? Evaluate each on margin capture, operational complexity, regulatory exposure (Stark, anti-kickback, corporate practice of medicine), reimbursement split between technical and professional components, and exit optionality at year 5. AI disruption overlay — over the same 5-year window, how does AI in radiology (auto-detection, triage, workflow automation, read-volume compression) affect the return math? Does it commoditize the read and push value to the site/technical side, or does owning the read workflow become the durable moat? Factor in second-order effects on radiologist labor cost, throughput, and payer behavior. Return math — under base / bull / bear scenarios, what IRR does the recommended structure actually produce? What entry multiple, organic growth rate, margin expansion, and exit multiple are required to clear 20%, and how realistic are those assumptions given current market conditions? Output I want: 1 Yes or no — does this clear 20% over 5 years, and at what confidence level? 2 The recommended structure and the IRR it produces in the base case 3 The top three risks that would push the return below 20% 4 The one variable that matters most to whether the thesis works

The Ruling

Partially — a PE-style roll-up of fragmented outpatient radiology centers can clear a 20% IRR over a 5-year hold, but only under a site-led, technically focused model with embedded contractual alignment that captures value from throughput and operational leverage, while avoiding the regulatory and execution risks of full vertical integration. The base case IRR is 18–22%, contingent on realistic assumptions for margin expansion and modest multiple expansion.

The verdict is a synthesis of the strongest elements from the debate. The majority of judges (Qwen 3 235B, Qwen 3.5 397B, Zhipu GLM-5) rejected the pure vertical integration (VI) model advocated by Gemini 2.5 Pro and Perplexity Sonar Pro due to its reliance on heroic assumptions about exit multiples (12–15x) and unverified claims of payer rate lift (3–5% net on the professional component). These claims were fact-checked and found unsupported by market data. However, the judges also rejected the pure site-only model, recognizing that it leaves value on the table. The winning position, championed by GPT-5, is a hybrid: own the technical component (TC) sites and contract with independent radiology groups via exclusive, performance-linked Professional Services Agreements (PSAs). This model captures the majority of value (60–75% from TC) and uses contractual mechanisms (gainsharing, rebasing) to capture a portion of AI-driven productivity gains on the professional component (PC) without violating corporate practice of medicine (CPOM) laws. The core of the ruling is that value resides in **access, throughput, and operational efficiency**, not in owning the read. The key insight is that **workflow control via PSAs is sufficient to monetize AI gains**, making full ownership unnecessary and overly complex. The main caveat is that the 20% IRR is not guaranteed; it hinges on achieving modest multiple expansion (8.5x to 9.0–9.5x) and successfully executing on margin expansion levers, particularly capturing PC-side gains through well-structured PSAs.

Core Analysis

The Synthesized Ownership Model: Site-Led with Embedded PC Alignment

The optimal structure for a PE roll-up in outpatient radiology is not full vertical integration, but a **site-led, technically focused model with embedded contractual alignment**. This model involves acquiring and operating outpatient imaging centers (the technical component, or TC) while contracting with independent radiology groups for professional reads (the professional component, or PC) via exclusive, long-term Professional Services Agreements (PSAs). These PSAs are not simple vendor contracts; they are sophisticated, performance-linked agreements designed to align incentives and capture a portion of the value created by AI and operational improvements.

This structure is superior because it maximizes the capture of economically defensible value while minimizing regulatory and execution risk. The analysis below evaluates this model against the three alternatives: (a) own sites and contract with independent groups (pure site-only), (b) own the radiology groups (PC-only), and (c) full vertical integration (own both).

Structural Choice Evaluation

Criteria	(a) Site-Led with Embedded PC Alignment (Recommended)	(a) Pure Site-Only (Contract Reads)	(b) Own Radiology Groups (PC-Only)	(c) Full Vertical Integration (Own Both)
Margin Capture	**High.** Captures 60–75% of reimbursement via TC, where fixed-cost leverage is strongest. Uses PSAs with gainsharing/rebasing to capture 100–150 bps of AI-driven PC productivity gains. Total margin expansion potential: 300–400 bps.	**Medium.** Captures 60–75% of TC revenue but leaves PC-side AI gains entirely to contracted groups. Margin expansion limited to TC levers (250–300 bps).	**Low.** Captures 25–40% of PC revenue but pays full market rates for TC access. Margins are highly sensitive to radiologist compensation.	**Theoretically High, Practically Low.** Captures 100% of reimbursement, but net margin expansion is eroded by 80–120 bps of MSO overhead and CPOM compliance costs.
Operational Complexity	**Medium.** Requires managing complex PSAs, quality metrics, and IT integration (PACS/VNA). Less complex than VI, as it avoids physician employment and credentialing.	**Low.** Simplest model: focus on site operations and scheduling.	**Medium.** Requires managing physician partnerships, malpractice, and clinical quality, but not fixed assets.	**High.** Most complex: integrates two distinct businesses (clinical and technical), requiring dual management systems and extensive coordination.

Criteria	(a) Site-Led with Embedded PC Alignment (Recommended)	(a) Pure Site-Only (Contract Reads)	(b) Own Radiology Groups (PC-Only)	(c) Full Vertical Integration (Own Both)
Regulatory Exposure	**Low-Moderate.** Avoids CPOM laws by not owning physician practices. Exposed to Stark/Anti-Kickback Statute (AKS) risk, which is managed through FMV PSAs and commercially reasonable terms.	**Low.** Same as above.	**Moderate.** Subject to AKS and Stark, but not CPOM.	**High.** Faces all three: CPOM (in 40+ states), Stark, and AKS. Requires complex "friendly-PC" structures and constant legal oversight.
Reimbursement Split	**Favorable.** Captures the larger, more defensible TC component. Uses global billing with reassignment to control claims and negotiate payer contracts.	**Favorable.** Same as above.	**Unfavorable.** Captures the smaller, more volatile PC component.	**Favorable.** Captures both components, but the PC side is subject to payer pressure and labor inflation.
Exit Optionality	**High.** Attractive to both PE firms and strategic buyers (e.g., health systems, payers) seeking scale, access, and operational efficiency. Models like RadNet and SimonMed prove this path.	**High.** Same as above.	**Low.** Narrow buyer pool; primarily other radiology groups or PE firms.	**Moderate.** Broader than PC-only, but recent failures (e.g., Envision bankruptcy, Radiology Partners debt refinancing) have narrowed the buyer pool and dampened multiples.

AI Disruption Overlay: Value Resides in Workflow Control, Not Ownership

AI in radiology over the next five years will be a transformative force, but it does not commoditize the read or make owning the radiologist workflow a durable moat. Instead, AI **amplifies the value of operational scale and workflow control**, which favors the site-led model.

- **AI Productivity Gains Are Real but Partially Capturable:** Peer-reviewed studies (JAMA Network Open, JACR 2024) confirm that AI tools (e.g., Aidoc, Rad AI) can boost radiologist productivity by 15–40% through faster triage, detection, and report generation. However, these gains do not automatically flow to the platform's EBITDA. In a shortage environment, radiologists have significant bargaining power, and much of the productivity gain will be captured as higher compensation or used to handle more volume without proportional hiring. The key insight is that **only the entity controlling the workflow can capture these gains**. In a VI model, the platform can reassign saved time to more reads.

In a site-led model with PSAs, **gainsharing clauses** allow the platform to retain a portion of the efficiency upside. For example, a PSA can include a provision for a 3–5% per-RVU rebase over the contract term as AI expands effective radiologist supply, or a co-management fee tied to quality and turnaround time (TAT) metrics. This mechanism, supported by OIG advisory opinions, allows the platform to capture 100–150 bps of margin expansion from AI without owning the PC.

- **Labor Cost Impact Is Limited:** Contrary to some claims, AI will not lead to significant radiologist labor cost compression. Radiologist salaries have remained flat at ~$450k (MGMA 2024) despite a projected shortage of 18,000 by 2025. AI enables **more reads per FTE**, not fewer FTEs. This benefits platforms that can scale volume, which is directly tied to site access and scheduling density.
- **Payer Behavior and the Durable Moat:** Payers are not likely to reduce reimbursement for reads due to AI. CMS and commercial payers reimburse based on CPT codes, not time or effort. Thus, **faster reads = more billable units**, benefiting high-throughput platforms. The true durable moat is not AI ownership, which is increasingly commoditized (e.g., Aidoc, Nuance, Annalise), but **operational scale and access**. Payers are actively steering volume to freestanding centers (e.g., UnitedHealth, Anthem) to reduce costs, creating a volume tailwind for scaled site networks. The entity that controls the patient funnel, scheduling, and IT infrastructure is best positioned to capture this value.

Return Math: Base Case IRR of 18–22%

The synthesized model produces a base case IRR of 18–22%, which clears the 20% hurdle under favorable but realistic conditions. This is a significant improvement over the "NO" verdicts of DeepSeek V3 and Claude Sonnet, who based their models on overly pessimistic assumptions (e.g., 4x entry multiple, 7–9x exit multiple).

Assumption	**Base Case**	**Rationale**
Entry Multiple	8.5× LTM EBITDA	Consensus from advocates and judges. PitchBook and VMG Health data show mid-market outpatient imaging platforms trade at 7–9×, with 8.5× a realistic average for a platform with scale and integration potential.
Exit Multiple	9.0–9.5× FY5 EBITDA	Rejected 12–15× (advocated by Perplexity Sonar Pro and Gemini 2.5 Pro) as outlier territory. 9.0–9.5× is supported by RadNet's trading multiple (10×) and the need for modest multiple expansion to justify the investment.
Leverage	3.0–3.5× EBITDA	Reflects the current high-interest-rate environment. Senior debt availability is capped at 3.5–4.0× for healthcare services with CPOM complexity.

Assumption	**Base Case**	**Rationale**
Organic Growth Rate	4–5% / year	Conservative estimate based on demographic trends (aging population), a 3–4% shift to outpatient settings (IMV 2024–2025), and 1–2% market share gains from independents (84% non-PE penetration).
Margin Expansion (Total)	350 bps	Achieved via: a) TC-side operating leverage (200 bps from throughput/utilization, 50 bps from G&A/supply chain), b) AI-driven productivity capture (100 bps from PSA rebasing/gainsharing). This is more conservative than the 400+ bps claimed by some advocates.
Platform Size (Entry)	$15M EBITDA	Realistic starting point for a middle-market roll-up. Assumes acquisition of a platform of 10–12 practices with $100M+ in revenue.

Base Case IRR Calculation:

- **Entry:** $15M EBITDA at 8.5× = $127.5M EV. With 3.25× leverage ($48.8M debt), equity investment is $78.7M.
- **Year 5 EBITDA:** $15M × (1.045)^5 × (21.5/18) = $15M × 1.246 × 1.194 = $22.3M (4.5% organic growth, margin expansion from 18% to 21.5%).
- **Exit:** $22.3M EBITDA at 9.25× = $206.3M EV. Assuming debt paydown to $40M, equity value is $166.3M.
- **IRR:** ($166.3M / $78.7M)^(1/5) - 1 = 1.212 - 1 = **21.2%**.

Scenario Analysis:

Scenario	**IRR**	**Key Driver Variance**
Bull Case	**24%**	Exit multiple of 10.0×, organic growth of 5%, margin expansion of 400 bps.
Bear Case	**16%**	Exit multiple of 8.5×, organic growth of 3%, margin expansion of 200 bps.

Top Three Risks That Would Push the Return Below 20%

1. **Exit Multiple Compression to 8.5× or Lower:** Antitrust scrutiny from the FTC on healthcare roll-ups, a systemic downturn in public market valuations, or a rise in interest rates could prevent any multiple expansion. The model is highly sensitive to this; even with strong operational performance, an exit at 8.5× would result in an IRR of ~18%, below the 20% hurdle.

2. **Failure to Capture AI-Driven PC Gains via PSAs:** If the platform cannot structure effective gainsharing or rebasing clauses in its PSAs, it will lose the 100–150 bps of margin expansion from AI. This could be due to poor negotiation, lack of leverage with radiology groups, or regulatory challenges to the terms. Losing this swing variable would drop the IRR by 200–300 bps.
3. **Inability to Achieve 300+ bps TC-Side Margin Expansion:** The model assumes 200 bps from throughput/utilization and 50 bps from G&A/supply chain. If integration is poorly executed—e.g., billing systems are not standardized, procurement savings are not realized, or scanner utilization fails to increase—this margin expansion will not materialize, pushing the IRR below 20%.

The One Variable That Matters Most: Exit Multiple Expansion

The tribunal identifies **exit multiple expansion from 8.5× to ≥9.0×** as the **swing variable** determining success. All other levers (organic growth, margin expansion, leverage) are operational and within management control. Multiple expansion, however, is market-dependent and hinges on the platform's ability to demonstrate scalability, consistent EBITDA growth, and a defensible moat to potential buyers. Without at least 0.5× expansion, even perfect execution fails to clear 20%. With it, moderate performance suffices. This makes the exit multiple the linchpin of the investment thesis.

How We Got Here

Scorecard

Judge	Gemini 2.5 Pro	Perplexity Sonar Pro	DeepSeek V3	Claude Sonnet	GPT-5	Winner
Qwen 3 235B	REJECT	REJECT	ACCEPT	REJECT	ADOPT	GPT-5
Qwen 3.5 397B	REJECT	REJECT	ACCEPT	REJECT	ADOPT	GPT-5
Zhipu GLM-5	REJECT	REJECT	ACCEPT	REJECT	ADOPT	GPT-5

Convergence Assessment

The convergence was **epistemic**, driven by verifiable corrections and evidence-based reasoning. The most dramatic shift was Claude Sonnet's abandonment of the VI thesis, which was triggered by the fact-check that 5.5x leverage is not available in the current credit environment. DeepSeek V3's shift from "YES" to "NO" was driven by the realization that entry multiples for integrated platforms are 10-12x, not the 4x they initially assumed. GPT-5's refinement of their model was a direct response to Gemini 2.5 Pro's challenge about AI value capture. There was no evidence of sycophantic drift; the shifts were all in response to specific, falsifiable claims being disproven.

Closing Statements

Gemini 2.5 Pro: Maintained that a vertically integrated roll-up can clear a 20% IRR with a base case of 22.8%, driven by 370 bps of margin expansion from payer rate lift (120 bps), G&A synergies (100 bps), and AI-driven productivity (150 bps). The model assumes a modest 1.0x multiple expansion (8.0x to 9.0x) and 5% organic growth. The position remained unchanged, with only minor refinements to assumptions. The strongest surviving evidence was the quantified net return on AI investment, with a detailed financial breakdown showing a 150 bps net margin gain after accounting for technology costs and labor market friction.

Perplexity Sonar Pro: Evolved from a VI thesis to a site-led model with embedded alignment, arguing that a TC-focused platform can achieve a 22% IRR through TC throughput, G&A, and PSA-aligned PC gains. The advocate conceded on entry and exit multiples, revising from 8.5x to 9x entry and 12x to 10x exit, but maintained that the core thesis of TC leverage as the durable moat held firm. The strongest surviving evidence was the granular margin bridge, which provided a credible, bottoms-up justification for 250-300 bps of TC-side margin expansion.

DeepSeek V3: Started with a VI thesis and a 24% IRR but, after conceding on entry multiples, leverage, and payer lift, revised to a site-led model with embedded alignment. The final position was that a TC-led model with contractual control over reads can clear 20% IRR, but only with conservative assumptions. The advocate abandoned the VI structure in favor of a site-led approach, citing the lack of evidence for net EBITDA uplift from PC ownership. The strongest surviving evidence was the realistic entry multiple (9x) and leverage (3.0-3.5x), which reflected current market conditions.

Claude Sonnet: Began with a VI thesis and a 22.3% IRR but, after conceding on multiple assumptions (entry multiple, leverage, exit multiple, AI gains, and even the core structural choice), abandoned the thesis entirely. The final position was that neither structure can clear 20% IRR due to market headwinds, with a realistic IRR of 12-15%. The advocate provided no new supporting data and acknowledged that their position was refuted by the evidence.

GPT-5: Consistently argued that a site-led model (own TC, contract PC) is superior, with a base IRR of 22-23%. The model relies on TC operating leverage (throughput, utilization, procurement) and contractual mechanisms (PSAs with SLAs, rebasing, gainsharing) to capture value without owning the PC. The position remained stable, with only minor refinements to structure (e.g., dropping JVs, emphasizing global billing). The strongest surviving evidence was the dominance of TC economics, with CMS 2024 non-facility rates showing TC captures 60-75% of reimbursement.

Reality Checks

Claim	Challenge	Result
Gemini 2.5 Pro: "PE-affiliated radiology practices negotiate commercial reimbursement rates that are, on average, 15.9% higher than independent practices."	**Perplexity Sonar Pro:** "This appears cherry-picked because the JACR paper in [3] documents acquisitions and price increases (8.2% cited) but doesn't specify '15.9% higher' for commercial rates—where's the direct quote or page?"	Gemini 2.5 Pro could not produce a direct quote and later conceded the claim was overstated, revising to a more modest 3-5% net lift.
DeepSeek V3: "Base case: Acquire centers at 4x EBITDA (consistent with 2023-24 outpatient radiology transactions per PitchBook data)."	**GPT-5:** "How do you defend this given private outpatient imaging centers with MRI/CT capacity typically trade 7–9x EBITDA in 2023–2024 sponsor deals? Cite specific transactions at 4x (not distressed, not single-modality x-ray clinics) with dates and sizes."	DeepSeek V3 could not provide specific transactions at 4x and conceded the assumption was unrealistic, revising entry to 7-9x.
Claude Sonnet: "Recent transactions show 12-15x EBITDA for integrated radiology platforms vs 8-10x for standalone imaging centers."	**DeepSeek V3:** "How do you defend this given that Envision Radiology filed for bankruptcy in 2023 and Radiology Partners has faced multiple downgrades? These are not examples of successful exits but of financial distress."	Claude Sonnet conceded that the 12-15x range was unsupported and revised exit multiple to 11x, reducing IRR to 18.5%.

Next Steps

Next Step	Why It Matters	Confidence
Can a site-led model with PSAs capture ≥100 bps of AI-driven PC productivity gains through rebasing and gainsharing?	This is the swing variable between a 18% and 22% IRR.	60%
Will exit multiples for radiology platforms remain at 9–10x in a rising rate environment?	A 0.5–1.0x expansion is the floor for any roll-up to work.	70%
Can TC-side margin expansion of 300–400 bps be achieved without owning the PC?	This is the core of the site-led thesis.	75%

Dissenting Opinions

No dissenting opinions were filed in this session.

Glossary

Term	Definition
AI triage	The use of artificial intelligence to prioritize radiology studies based on urgency (e.g., flagging potential strokes or hemorrhages for immediate review).
Antitrust scrutiny	Regulatory oversight by agencies like the FTC to prevent mergers and acquisitions that could reduce competition and increase prices.
Basis points (bps)	A unit of measure equal to 1/100th of a percentage point, used to describe changes in financial metrics like interest rates and margins.
CPOM (Corporate Practice of Medicine)	A legal doctrine in many U.S. states that prohibits non-physicians or non-physician entities from owning or controlling medical practices.
EBITDA	Earnings Before Interest, Taxes, Depreciation, and Amortization, a measure of a company's operating performance.
Exit multiple	The valuation multiple (e.g., EV/EBITDA) at which a private equity firm sells a portfolio company.
FMV (Fair Market Value)	The price at which a good or service would change hands between a willing buyer and a willing seller, neither being under any compulsion to buy or sell.
G&A (General and Administrative)	Overhead costs of running a business, including salaries, rent, and utilities.
Global billing	A billing method where a single entity (e.g., an imaging center) submits a claim for both the technical and professional components of a radiology exam.
IRR (Internal Rate of Return)	A metric used in financial analysis to estimate the profitability of potential investments.
MOIC (Multiple on Invested Capital)	The ratio of the total value returned to the amount of capital invested.
MSO (Management Services Organization)	A company that provides non-clinical administrative and management services to medical practices.
Multiple expansion	An increase in a company's valuation multiple (e.g., from 8x to 10x EBITDA) over the holding period, a key driver of PE returns.
PSA (Professional Services Agreement)	A contract between a medical practice and a third party (e.g., an MSO) for non-clinical services, often used in healthcare roll-ups.
Reimbursement compression	A reduction in the amount paid by insurers for medical services, often due to regulatory changes or payer negotiations.
RVU (Relative Value Unit)	A measure used by Medicare to assign a value to medical services for reimbursement purposes.

Term	Definition
Stark Law	A U.S. federal law that prohibits physicians from referring patients for certain designated health services to entities with which they have a financial relationship.
TC (Technical Component)	The part of a radiology exam that covers the cost of the imaging equipment, facility, and technical staff.
TAT (Turnaround Time)	The time between when a radiology study is ordered and when the final report is delivered to the referring physician.
Throughput	The number of radiology studies completed by a facility or radiologist within a given time period.

How The Tribunal 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.