BUILDING AI-NATIVE PROFESSIONAL SERVICES FIRMS

Strategy, Economics, and Execution

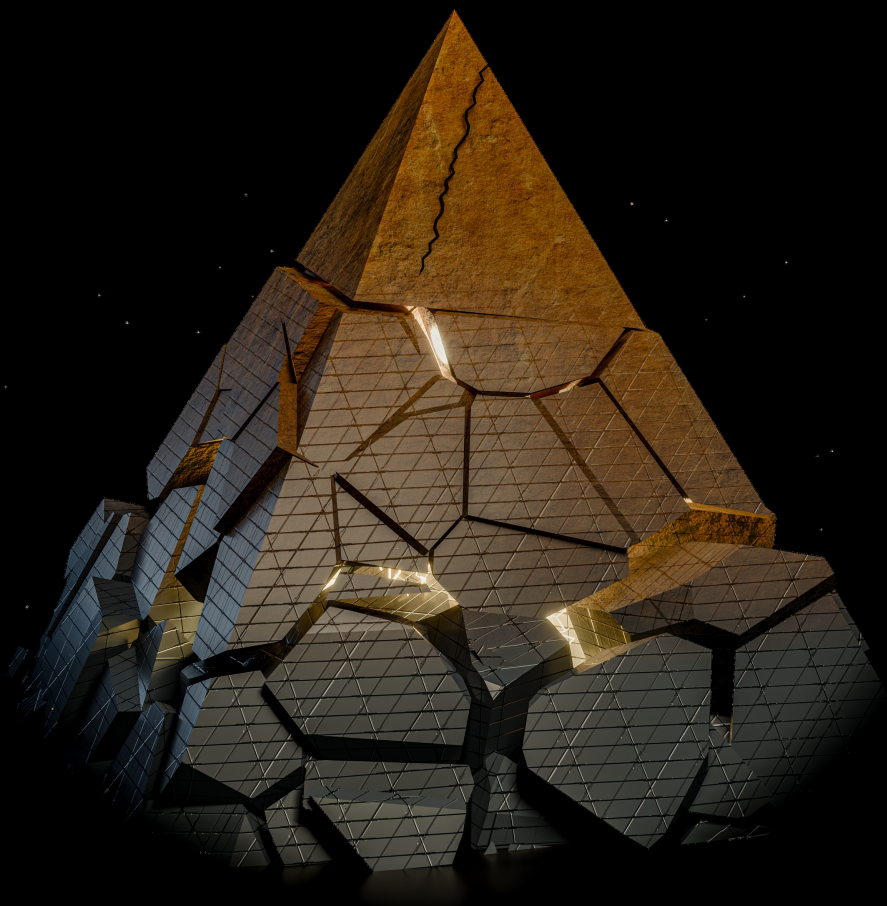

DANIEL M. KATZ · MICHAEL J. BOMMARITO II · JILLIAN BOMMARITO

Chapter 2

The Economics of Transformation

Chapter 2 · March 2026

Chapter 2: The Economics of Transformation

All models are wrong, but some are useful.

— George E. P. Box, 1976

Sarah arrived at the restaurant fifteen minutes early. Maya had picked the place—a farm-to-table spot in the West Village where every dish came with a paragraph on the provenance of its ingredients, the kind of restaurant where New York dealmakers met when they wanted to signal taste rather than just money. The menu changed daily and the wine list ran to forty pages. Sarah chose a table in the back corner—quiet enough for the conversation she needed to have, visible enough that Maya would spot her when she arrived.

She had not seen Maya Chen since a brief coffee two years ago, though they texted occasionally. They had met freshman year at an elite East Coast university where half the students seemed destined for Wall Street and the other half for elite law firms. She and Maya were both pre-law and ended up in the same microeconomics section. Maya had pivoted to finance after graduation while Sarah went to law school. Now Maya was quickly climbing the ranks in the Professional Services M&A practice at one of the bulge bracket investment banks. It was one of those fortunate accidents of life—that a college friend had gone on to develop deep expertise in exactly the domain Sarah now needed to understand. Maya had spent years advising on mergers, acquisitions, and capital raises for professional services firms. She knew how they were valued, how they

were structured, and where they broke. Not many people on earth had that particular combination of knowledge and willingness to share it over dinner.

Sarah was four months into her new firm. She had two associates, a handful of clients, and a growing conviction that she was onto something real. But conviction was not enough. She needed to understand the economics—really understand them—both for her own planning and for the investor conversations that were starting to materialize. Maya had valued dozens of professional services firms. She had seen what worked and what did not. If anyone could give Sarah the inside view on how the financial world actually thought about firms like hers, it was Maya.

Sarah had done the math more times than she cared to admit. She had been able to bring some clients and revenue with her. Even if that dried up, she had a runway of six months, maybe twelve. If she could land a few more clients she could stretch it further. But the opportunity costs were significant. She had left a $385,000 salary and a clear path to partnership—the kind of path her law school classmates would have fought for. Her former mentor had called it reckless. Her parents had called it brave, which Sarah suspected was their polite word for the same thing.

Maya appeared at 7:15, still in her work clothes, phone pressed to her ear. She held up one finger—almost done—and slid into the booth across from Sarah.

"Sorry, sorry. Deal closing. You know how it is." She dropped her phone face-down on the table. "Actually, you probably do not know how it is. You escaped this life."

"I escaped to build a law firm. Which may be worse."

Maya laughed. "I heard. An AI-native law firm. I have been seeing a lot of those pitches lately." She picked up the wine list. "Tell me why yours is different?" Maya was a friend but she was all about business. And questions like that were exactly why Sarah had made the trip.

2.1 Enterprise Value, Decoded

"Okay. Let us start with the question every managing partner asks me, usually with a mix of hope and terror: What is my firm worth?"

"And what do you tell them?"

"I tell them it depends on what someone is willing to pay. Which sounds like a non-answer, but it is actually the only honest answer."

"But there must be a more rigorous way to think about it."

"There is. It is called enterprise value." Maya pulled a pen from her bag and flipped over her placemat. "Let me show you the framework we actually use."

She sketched on the paper:

Enterprise Value

Enterprise value is the present value of expected future cash flows, plus a terminal value representing the firm's worth beyond the projection period. Future cash flows are discounted at a rate that reflects the risk of the business.

Sarah studied the placemat. "I remember some of this vaguely from undergrad. You know, I went into law partly to get away from all of this." She smiled. "Guess the joke is on me."

"Right. And here is my confession before we go further." Maya set down her pen. "This framework is both profoundly useful and deeply problematic. It requires predicting future cash flows that no one can predict with confidence. It collapses all the uncertainty facing a business into a single discount rate. It rests on assumptions about perpetual growth that strain credibility for any business, let alone one undergoing transformation."

"Then why use it?"

"Because it is the common language of capital allocation. When you negotiate with investors, acquire a competitor, or evaluate a strategic decision, the other

party will speak this language. You need to speak it too—even as you understand its limits."

Sarah nodded slowly. "Well, that is why I am here. I need a crash course, so please—teach me to speak the lingua franca."

Maya ordered appetizers and continued. "Let's just take this one metric at a time."

EBITDA

EBITDA—Earnings Before Interest, Taxes, Depreciation, and Amortization—measures a firm's operating profitability by stripping out financing decisions, tax strategies, and accounting conventions. It is the standard valuation metric in professional services transactions, where enterprise value is typically expressed as a multiple of EBITDA.

"EBITDA is the operating engine of the firm, stripped of everything that is just accounting or capital structure. It is the number that gets multiplied in every professional services transaction. When I evaluate a firm, revenue tells me how big it is. EBITDA tells me how good it is."

"So a firm could have growing revenue but terrible EBITDA," Sarah said.

"Happens all the time. And for years—the entire ZIRP era, really, when interest rates were essentially zero—growth was all anyone seemed to care about. Cheap capital everywhere. Investors would fund companies burning cash at staggering rates as long as the growth curve pointed up and to the right." Maya shook her head. "Then rates went up and the music stopped. Capital is expensive now, and investors want to see a path to profitability. EBITDA has become one of the key metrics that separate real businesses from stories."

"But EBITDA alone is not enough," Maya continued. "You also need to understand what is actually left in the bank."

Free Cash Flow

Free cash flow is what remains after the firm generates earnings from operations, pays taxes, invests in long-term assets such as technology platforms, and funds the timing gap between delivering work and collecting payment.

"EBITDA and free cash flow answer different questions. EBITDA tells an investor what the business is worth. Free cash flow tells you whether the business actually works—whether you can fund growth, service debt, and distribute to owners. A firm with strong EBITDA but negative free cash flow looks valuable on paper but cannot sustain itself. Law firms know the working capital problem well—you bill in October, the client pays in January, and for three months that cash is trapped in receivables."

Sarah paused. "At my old firm, the partners never discussed either one. It was always revenue and billable hours. They tracked the top line obsessively but never asked what was left after the firm paid its operating costs, invested in the business, and funded working capital."

"That is the blind spot," Maya said. "Most law firm partners have never had to think like business operators. The traditional model generated enough profit that financial sophistication was often optional."

"Let me give you an analogy," Maya said. "Imagine two restaurants. Both do $2 million in revenue. Restaurant A spends $1.8 million running the place—high rent, expensive staff, spoilage, chaos in the kitchen. EBITDA: $200,000. Restaurant B has the same revenue but a tight operation—smart sourcing, less waste, technology in the back of house. EBITDA: $600,000. A buyer shows up. They are not paying for the revenue. They are paying for the profit engine. Restaurant B sells for three times what Restaurant A sells for—on the same top line."

"And the AI-native firm is Restaurant B," Sarah said.

"The AI-native firm is Restaurant B with the potential to become a chain. That is the multiple expansion. You are not just more profitable—you are more scalable. And scalability is what separates a 1–2x revenue multiple from a 5–10x

multiple in the eyes of an investor." Maya leaned forward. "But here is the part that most people miss. Restaurant A cannot become Restaurant B by buying a new oven. You have to redesign the kitchen. Change the menu. Retrain the staff. Rethink what you are actually selling. That is what AI-native means—it is not a tool you bolt on. It is a fundamentally different operating model. And the market will price that difference accordingly, but only if the EBITDA proves it is real."

Sarah sat with that for a moment. "So the EBITDA is the proof. The revenue is just the stage."

"Now you are thinking like an investor," Maya said.

"One more piece," Maya said. "Remember that discount rate in the enterprise value formula? Investors discount future cash flows based on risk. The riskier the business, the less those future earnings are worth today. Traditional firms carry heavy risk—clients follow partners, revenue tracks headcount, client switching costs are low. Those factors depress valuations." She paused. "But here is the key insight: small reductions in risk create outsized increases in value. Embedding knowledge in systems instead of people, diversifying client concentration, building defensible moats—these moves can create more enterprise value than margin improvement alone."

Sarah leaned back. "So risk reduction can create more value than margin improvement."

"A firm that reduces its risk profile while expanding margins is pulling two levers at once—and the risk lever is the more powerful of the two."

2.2 Growth Without EBITDA: A Cautionary Tale

Maya paused. "Let me tell you a story that makes this concrete. Have you heard of DWF?"

"The UK firm? I know the name."

"DWF was a top-30 UK law firm that went public on the London Stock Exchange in March 2019 (DWF Group plc 2019). Premium listing, Main Market—not some junior exchange. They raised £75 million at 122 pence per share, giving them a market capitalization of £366 million. It was the largest listed law firm in the world."

"Was," Sarah said.

"Was." Maya set down her glass. "Here is what happened. To make the economics work for external shareholders, DWF equity partners accepted a 60 percent reduction in their annual compensation (Legal Business 2019). They traded current income for shares. The firm then spent aggressively on acquisitions—a Spanish firm for over 40 million euros, a legal process outsourcing business in India for $18.5 million—funded largely with debt. Revenue grew from £272 million to over £450 million in four years."

"That sounds like a success story."

"Look at the EBITDA. That is the point." Maya tapped the table for emphasis. "Revenue grew 66 percent. But profitability never followed. When COVID hit in 2020, DWF issued a profit warning and scrambled to arrange emergency credit facilities. The share price collapsed. The CEO who had championed the IPO was ousted 14 months after listing—the stock dropped 18 percent in a single day (City A.M. 2020). At its lowest point, shares traded at 53 pence. Less than half the IPO price."

Sarah winced. "What happened to the partners who took the pay cut?"

"They were locked into a five-year vesting schedule, watching their shares lose value while earning 60 percent less than competitors at traditional firms. DWF was trying to recruit talent against firms distributing 100 percent of profits to partners." Maya shook her head. "In October 2023, a private equity firm bought DWF for 100 pence per share—an 18 percent discount to the IPO price (Inflexion Private Equity 2023). Four and a half years as a public company, and investors lost money. The firm spent £22 million on IPO expenses alone—nearly 30 percent of gross proceeds."

"So the public market experiment failed."

"For DWF, yes. And the irony is that the firm has reportedly performed much better since going private (DWF Group 2024). Revenue is growing, profitability has improved. The lesson is not that DWF was a bad firm. The lesson is that public markets are unforgiving for professional services businesses with thin margins and lumpy revenue. One bad quarter triggers a profit warning, which triggers a share price collapse, which triggers a leadership crisis." Maya met Sarah's eyes. "Revenue told the market DWF was getting bigger. EBITDA told the market it was not getting better. And with respect to professional services firms, the market largely cares about the second number."

Sarah sat with that for a moment. "So when you said EBITDA separates real businesses from stories—"

"DWF had a compelling story. Growing revenue, global expansion, the prestige of being the only law firm on the LSE's premium segment. But the underlying economics—the EBITDA, the margins, the free cash flow—never supported the valuation the market assigned at IPO."

"What multiple did DWF actually trade at?"

"1.55 times revenue. £366 million market cap on £236 million of pre-IPO revenue." Maya let the number hang in the air. "And that is instructive, because it tells you exactly how the market prices a traditional professional services firm—even one with a growth story and a prestigious listing. 1.55 times revenue. Compare that to a mature software company, which typically trades at 5 to 10 times revenue. A high-growth SaaS business with strong retention metrics might command 10 to 15 times."

Sarah leaned forward. "And what about AI companies?"

Maya exhaled. "This is where things get interesting—and, frankly, a little unhinged. The median late-stage AI startup trades at roughly 25 times revenue. The outliers are much higher. OpenAI has been valued at roughly 60 times revenue. Some AI companies have raised capital at over 100 times revenue."

"And legal AI specifically?"

"Legal AI is in its own category right now." Maya pulled up something on her phone. "I have seen legal AI companies—backed by the biggest names in venture capital—reach valuations of $5 to $8 billion on annualized revenue of

$150 to $200 million. That is a multiple of 40 to 50 times revenue. Some smaller platforms have hit $1 to $2 billion valuations on revenue that might be $30 to $40 million. Roughly 45 times."

Sarah blinked. "Forty-five times revenue. DWF got 1.55 times."

"There are a few exceptions—legal technology companies that have been around longer, built real EBITDA, and trade at a more reasonable 10 to 15 times revenue on hundreds of millions in recurring revenue. Those are defensible businesses with compounding customer data." Maya set down her phone. "But they are the exceptions. Most of the legal AI companies commanding these valuations? I would be very cautious."

"Why?"

2.2.1 *Don't Go Chasing Waterfalls*

"Because these are venture valuations, and venture valuations come with venture capital structures. Liquidation preferences. Participation rights. Anti-dilution protections." Maya ticked them off on her fingers. "You are a lawyer, I am sure you know these things. When a company raises a billion dollars at a $5 billion valuation, the investors are not simply betting that the company is worth $5 billion today. They are buying preferred shares with downside protection built in. If the company exits for less than the last-round valuation, the preferred shareholders get paid first—sometimes multiple times their investment—before a single dollar flows to founders, employees, or common shareholders."

"The waterfall."

"Exactly. And here is what I see when I look at legal AI." Maya lowered her voice. "Many of these companies will not grow into their valuations. Revenue multiples of 40 to 50 times require sustained hypergrowth—doubling or tripling revenue year after year—in a market where enterprise legal departments are cautious buyers, regulatory constraints limit adoption speed, and there is a somewhat limited moat relative to other competitors. When growth decelerates, and it will, these companies face down rounds. New investors come in at lower valuations with even more protective terms. Each round dilutes the founders

and early employees further. By the time a company exits—if it exits—the liquidation preferences and participation rights can consume most or all of the proceeds."

"So the founders end up with nothing?"

"Not nothing, necessarily. But far less than the headline valuation suggests. I have seen deals where a company exits at $500 million and the founders walk away with single-digit millions after the preference stack unwinds. The venture investors are protected. The founders and employees who took below-market salaries for years are not."

Sarah thought about the investor conversations she had been having. The term sheets she had not yet seen. "So what is the right approach?"

"No matter what they tell you, raise what you need, not what you can. Grow into your valuation rather than ahead of it. And never forget that a 10 times multiple on a profitable, capital-efficient business is worth more to you personally than a 50 times multiple on a cash-burning business with a preference stack three layers deep." Maya smiled. "DWF teaches you what happens when a traditional firm overreaches into public markets. The legal AI bubble will teach a different lesson—what happens when venture hype detaches from economic reality. Both end the same way. Someone does not get paid what they expected."

2.3 The Economics of Time-for-Money

The wine arrived. Maya poured two glasses and shifted her approach.

"Before explaining to an investor why AI-native is different, I think you need them to understand what you are disrupting. In other words, walk them through the economics at places like your old firm."

Sarah thought back. "We had about 120 lawyers. Revenue around $88 million. Partners billed at over a thousand an hour, associates starting at $395. The firm had been flat for three years, but the partners still expected their draws to increase."

"That is the staffing pyramid in action." Maya wrote on the placemat:

Traditional Revenue Model

$$Revenue = Hours \times Rate \times Realization \times Utilization$$

"Most associates and partners have goals in terms of hours billed, so target hours are the raw input. Rate is the price per hour. Realization is the percentage of billed time that clients actually pay—write-downs, disputes, discounts. Utilization is the percentage of available time that gets billed."

"And the model only works if you can keep pushing those numbers up," Sarah said.

"Right. But all four have practical ceilings. Hours are limited by available time. Rates face competitive pressure. Realization reflects client willingness to pay. Utilization reflects demand and efficiency." Maya sipped her wine. "The industry benchmark for utilization sits at 68.9 percent—well below the 75 percent target most firms need for healthy economics."

"My old firm was at 71 percent. They celebrated it."

"They should not have."

2.3.1 *The Cost Structure*

"The dominant cost is people. At most firms, professional and support compensation consumes 60 to 75 percent of revenue. This creates the economic incentive for staffing pyramids—if junior professionals generate revenue at rates higher than their fully-loaded cost, the firm captures the spread."

Sarah did the math in her head. "So if an associate bills at $400 an hour and costs $200 fully loaded, each hour of associate time generates $200 in margin."

"Exactly. If a partner supervises four associates billing 1,600 hours each, the partner's practice generates $1.28 million in margin from associate labor alone—before the partner's own billing." Maya set down her phone. "This model has generated extraordinary wealth. Partners at elite firms earn $3 to $10 million

annually. The Am Law 100 reported average profits per equity partner of $3.15 million in 2024."

"But it is under stress."

"Significant stress. Many Am Law 100 firms now pay first-year associates $225,000 in salary plus bonus. Senior associates at elite firms earn $400,000 or more. The margin on associate labor has compressed." Maya paused. "And the staffing ratios are shifting. Non-equity partners now represent 51 percent of the Am Law 100 partnership—up from 46 percent a decade ago."

"The pyramid is flattening."

"Whether they acknowledge it or not."

Sarah frowned. "But Big Law still dominates, doesn't it? The profits per partner numbers are staggering."

"The profits are real, but the market share is not what it used to be. Big Law has been losing ground for years. Outsourcing to alternative legal service providers and legal process outsourcing firms has grown steadily. Corporate legal departments have insourced aggressively—in-house headcount has expanded for two decades running. Technology and automation handle work that once required junior associates. These are not new trends. They have been compounding quietly." Maya set down her pen. "What has masked the erosion is that regulation and regulatory complexity, along with globalization, have generated enough high-value problems to keep Big Law churning. The total addressable market grew fast enough to offset the share losses. Whether that continues is not entirely clear, but there are cracks forming—particularly at certain firms that built their models around volume work that is now migrating elsewhere."

Maya pointed her pen at Sarah. "Now, I could spend all night talking about the structural problems facing traditional firms. But that is not why you are here. The only question that matters for you is whether there is room to build something viable and profitable inside this ecosystem. Let the incumbents worry about their pyramids."

2.3.2 *Why Multiples Are Low*

Sarah asked the question she had been building toward. "So what does all this mean for valuation?"

"Traditional professional services firms trade at modest multiples. Typically 1 to 3 times revenue, or 6 to 12 times EBITDA." Maya counted on her fingers. "Four factors explain this ceiling."

"Key-person risk comes first. Clients often follow partners, not firms. When a partner leaves, revenue follows. I advised on a deal last year where the target lost 22 percent of revenue within months of closing because two partners lateraled to a competitor."

"Limited scalability compounds the problem. Revenue grows only as fast as headcount. Unlike software companies with near-zero marginal costs, professional services firms must hire to grow. Every dollar of new revenue requires a proportional investment in people."

"Low switching costs make matters worse. Clients can change providers relatively easily. There is no lock-in, no network effect, no proprietary technology that makes switching expensive."

"Beyond the significant regulatory limitations, the partnership structure creates the final constraint. Traditional partnerships distribute most profits each year rather than retaining earnings for growth. This limits capital for investment and makes acquisition by outsiders complicated."

Sarah thought about her old firm's partner distributions. The pressure to maximize current draws rather than invest in the future. "So a well-run traditional firm might command 1.2 times revenue."

"An exceptional firm with growth momentum might reach 2 times. Software companies with similar revenue routinely trade at 5 to 10 times or more." Maya leaned forward. "That gap between professional services multiples and technology multiples is the opportunity you are chasing."

2.4 Articulating the AI-Native Advantage

Their entrees arrived. Sarah barely noticed.

Maya continued to provide her perspective. "AI and the future of professional services—this is where I have been spending most of my time lately. I have seen pitch decks from a range of so-called AI-native law firms, accounting firms, management consulting, financial advisory, and others. It is easy to use the term, but it is much harder to make it so. Some of it is hype. Some of it is real."

"I have my own ideas in mind, but what in your perspective makes it real?" asked Sarah.

Maya was quickly running out of cocktail napkins, but the waiter graciously brought her a small stack. "Whether it is a law firm, accounting firm, or some other professional services organization run as a partnership, AI-native firms directly address each of the four constraints I mentioned earlier, all of which historically depress traditional valuations: key-person risk, limited scalability, low switching costs, and partnership structure. If executed correctly the transformation is not incremental—it restructures the fundamental drivers of value."

2.4.1 *Technology Productivity Meets the Labor Pyramid*

"In the traditional model, the staffing pyramid means junior humans doing work for senior humans. In the AI-native model, AI systems handle a significant number of tasks that previously went to junior associates."

Maya wrote some numbers. "The arithmetic is stark. An AI system processing contracts, supporting regulatory requests or undertaking a range of similar law-related tasks does not earn $225,000 per year. It does not demand bonuses, take vacations, or leave for competitors. The cost of an AI worker is the amortized cost of developing the system plus marginal costs for compute—a fraction of human labor cost."

"What kind of productivity gains are you seeing?"

"Early data from AI-native firms suggests substantial shifts. Some claim lawyers using their platforms can process contracts 5 to 10 times faster than traditional approaches. Others argue that one lawyer can do the work of many

with proper AI augmentation." Maya held up a cautionary hand. "These are claims made by company founders that require independent verification. Even conservative estimates suggest 2 to 3 times productivity improvements at the task level, though engagement-level improvements are typically more modest."

Sarah shook her head. "I have seen it myself, but it is more nuanced than the pitch decks suggest. The contract review that took me eight hours now takes roughly four with AI assistance—but that is the structured, predictable work. The judgment calls, the negotiation strategy, the part where you have to tell a client something they do not want to hear—AI does not touch that. The founders claiming 10x are measuring the easy part."

Maya raised an eyebrow. "That is the most honest thing I have heard from anyone building in this space."

"I agree but notwithstanding, I do think this changes the staffing calculation. Instead of a partner supervising four associates, a partner might supervise four AI agents with one associate for quality verification. The profit on each unit of output expands because labor costs have dropped substantially."

2.4.2 *The New Cost Structure*

"Look at the key shifts," Maya said. "Professional compensation drops from 60 percent to 30 percent of revenue. Fewer humans generate the same or more output. Technology costs might rise from 3 percent to 20 percent—but that is still far less than the labor cost you saved."

The restaurant had filled up around them. A birthday party at the long table erupted in applause, but neither of them looked up.

Sarah studied the comparison. "The margin expansion is dramatic."

"From 10 percent to 36 percent. And that is before you factor in scalability."

2.4.3 *Scalability Changes Everything*

"As I noted earlier, traditional professional services have high marginal costs," Maya continued. "Serving one more client requires hiring more people. Growth requires proportional investment in headcount."

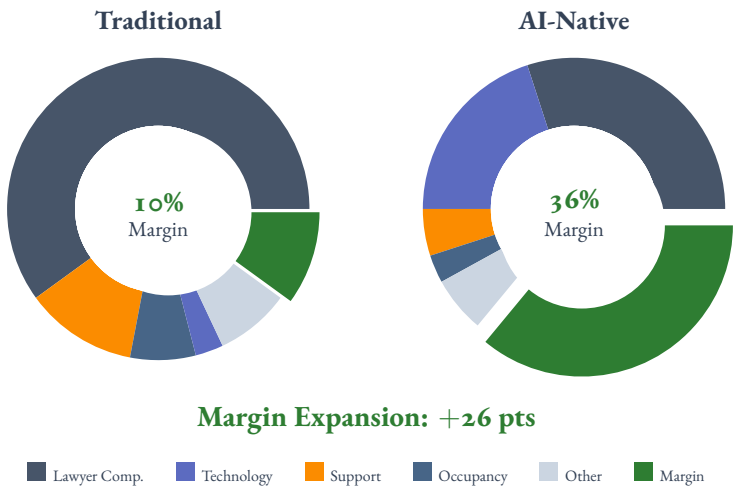

Figure 2.1: Cost structure shift: traditional firms spend 60% on lawyer compensation with 10% margins; AI-native firms substitute technology for labor, achieving 36% margins.

"But AI-native firms approach software economics," Sarah said, seeing where this was going.

"Exactly. Once the platform is built, the marginal cost of serving additional clients is primarily compute cost—which is low and declining. A firm that can serve 100 clients can serve 200 clients with modest incremental investment." Maya paused. "A traditional firm growing 20 percent per year must hire 20 percent more people. An AI-native firm growing 20 percent per year might hire 5 percent more people while investing in platform improvements."

2.4.4 *Multiple Expansion*

"So the economics are better," Sarah said. "But you mentioned something earlier about multiple expansion. Why would AI-native firms command higher valuation multiples?"

"I know I sound like a broken record but these are the points you need to drive home with investors over and over." Maya once again enumerated the

reasons. "Reduced key-person risk—knowledge is embedded in systems, not people. When a senior professional leaves, the platform retains the institutional knowledge that person contributed."

"Increased scalability—investors value growth, and AI-native firms can grow faster with less capital."

"Stronger moat—proprietary AI systems and accumulated data create barriers to entry that traditional professional services firms lack. Data advantages compound over time."

"Better unit economics—higher margins and lower customer acquisition costs improve the lifetime value to customer acquisition cost ratio that drives valuation in growth businesses."

Sarah thought about the investor conversations she had been having. "What kind of multiples are you seeing?"

"Early evidence is suggestive but not conclusive. Several AI-native legal platforms and firms have reached significant valuations—albeit venture valuations." Maya set down her glass. "The multiples on these companies—5 times revenue or higher—are unheard of in traditional professional services."

"But?"

"But venture valuations are not exit multiples. These are pre-revenue or early-revenue companies valued on potential, not proven economics. We have not yet seen an AI-native professional services firm exit at these multiples." Maya's expression turned serious. "Traditional professional services M&A typically commands 1 to 2 times revenue. ALSPs have occasionally achieved 2 to 3 times revenue. Whether AI-native firms can sustain technology-company multiples at scale remains unproven."

Sarah pushed her plate aside. "Can you show me the math? A direct comparison."

Maya pulled out her phone and opened a spreadsheet. "Let me walk you through a hypothetical. Consider a mid-market law firm with $50 million in revenue."

2.4.5 *Traditional Firm P&L*

Line Item	Traditional	%
Revenue	$50,000,000	100%
Professional compensation	$30,000,000	60%
Support staff	$6,000,000	12%
Occupancy	$3,500,000	7%
Technology	$1,500,000	3%
Other overhead	$4,000,000	8%
EBITDA	**$5,000,000**	**10%**

Table 2.1: Traditional Law Firm P&L ($50M Revenue)

"This firm employs approximately 50 lawyers at $600,000 average revenue per lawyer, plus 30 support staff. Partner profits come from the EBITDA line plus partner compensation embedded in professional costs." Maya scrolled. "Enterprise value at 1.2 times revenue: approximately $60 million."

2.4.6 *AI-Native Firm P&L*

"Now consider an AI-native firm serving similar clients with similar revenue."

"This firm employs approximately 18 lawyers—10 senior and 8 junior—plus 8 engineers and 10 support staff. Revenue per lawyer is roughly $2.8 million. The engineers are not an expense line you can cut—they are the reason the economics work." Maya looked up. "Enterprise value at 5 times revenue: approximately $250 million."

Sarah stared at the numbers. She pulled the phone closer, tilting it against the candlelight. "That cannot be right."

"Run the numbers yourself."

Line Item	AI-Native	%
Revenue	$50,000,000	100%
Professional compensation	$15,000,000	30%
AI/Technology	$10,000,000	20%
Support staff	$2,500,000	5%
Occupancy	$1,500,000	3%
Other overhead	$3,000,000	6%
EBITDA	**$18,000,000**	**36%**

Table 2.2: AI-Native Law Firm P&L ($50M Revenue)

Sarah studied the two columns. "But the margin expansion is not just about replacing associates with AI. Your AI-native P&L shows professional compensation at 30 percent of revenue instead of 60. That means the lawyers who remain—even the senior ones—are earning less in annual compensation than they would at a traditional firm."

"Correct. And that is by design." Maya pointed to the AI-native column. "In a traditional partnership, virtually all profit flows out as annual partner draws. There is nothing left to compound. In an AI-native structure with outside capital, lawyers accept lower current compensation in exchange for equity—ownership stakes whose value grows as the enterprise grows. A partner at a traditional firm might earn more than $1 million a year and yet own nothing transferable. A senior lawyer at an AI-native firm might earn half that amount but hold equity worth multiples of that difference when the firm reaches scale."

"So you are trading current income for terminal value."

"Exactly. The margin expansion is real, but part of it comes from professionals accepting a different compensation bargain—less cash today, more ownership value tomorrow. That only works if the enterprise actually builds value. Which is why the economics have to be right, not just the technology."

2.4.7 *The Value Creation Gap*

Maya showed her a summary table.

Metric	Traditional	AI-Native	Difference
Revenue	$50M	$50M	—
EBITDA	$5M	$18M	+$13M (260% increase)
EBITDA margin	10%	36%	+26 pts
Revenue multiple	1.2x	5.0x	+3.8x
Enterprise value	$60M	$250M	+$190M (317% increase)

"The value creation comes from two sources," Maya explained. "First, margin expansion increases EBITDA from $5 million to $18 million—a 260 percent improvement. Second, multiple expansion increases the valuation multiple from 1.2 times to 5 times—reflecting the firm's transformed risk profile and growth potential."

Sarah sat back. "$190 million in enterprise value. From the same revenue base."

"That is the prize," Maya said. "Firms that successfully transform can create extraordinary value. Sponsors who invest in transforming firms can capture significant returns."

"You said successfully," Sarah noted. "What about firms that do not succeed?"

2.5 The J-Curve and the Politics of Change

Maya's expression shifted. "This is the part that kills most transformations."

She drew a curve on a fresh napkin.

"The comparison I just showed you displays the before and after states—the equilibrium conditions years apart. But firms do not teleport from one state to the other. What happens during the transformation itself?"

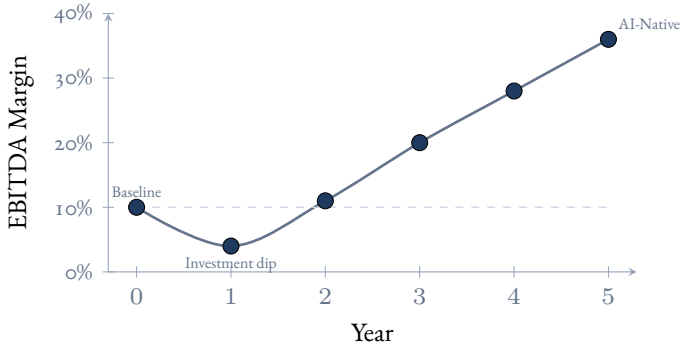

Figure 2.2: The transformation J-curve. EBITDA margin drops in Year 1 as technology investment precedes headcount reduction, then expands as AI productivity gains compound.

"The economics typically follow a J-curve," Maya said. "Margins initially compress before expanding."

She walked Sarah through the trajectory.

"Year 0 is your baseline. The firm has $50 million revenue with a 10 percent EBITDA margin producing $5 million EBITDA. Headcount is 50 lawyers and 30 support staff. Technology spend runs $1.5 million."

"Year 1 is the investment phase. Revenue is flat or modest growth at $52 million. Technology investment jumps to $8 million for platform development, AI tools, and training. Headcount remains unchanged initially—you cannot reduce staff until AI is proven. EBITDA compresses to $2 million. A 4 percent margin."

Sarah felt a knot in her stomach. "That is a 60 percent drop in profits."

A waiter appeared to clear their plates. Sarah waved him off without looking up.

"Welcome to transformation." Maya continued. "Year 2 brings early returns. Revenue grows to $55 million as efficiency enables more capacity. Technology stabilizes at $9 million. Professional headcount reduces by 15 percent to 43 lawyers. EBITDA recovers to $6 million—an 11 percent margin."

"Year 3 is acceleration. Revenue reaches $60 million. Technology holds at $10 million. Professional headcount drops to 30 lawyers—40 percent reduction from baseline. EBITDA reaches $12 million, a 20 percent margin."

"Year 5 is mature AI-native. Revenue has grown to $75 million, enabled by increased capacity. Technology runs $15 million. The firm operates with just 20 senior lawyers. EBITDA reaches $27 million—36 percent margin."

Sarah did the math. "So you have to survive a 60 percent profit drop before you see the gains."

"It does not have to be that dramatic. This is merely an illustration but the point is that this is where most transformations die." Maya set down her pen. "Do you have enough runway to get through the dip?"

Sarah thought about her firm's finances. Three people. A handful of clients. No legacy costs but also no cushion. "I am building AI-native from scratch. I do not have the dip problem."

"No. You have a different problem. You have to build the platform and acquire clients simultaneously, without the revenue base a traditional firm starts with." Maya paused. "But you also do not have partners voting to preserve their distributions."

"Which means I am going to need capital," Sarah said. "Real capital. Not just my savings and a line of credit."

"Exactly. You are looking at a seed round to get the platform built and the first clients onboarded, and then—if the unit economics prove out—a Series A to scale." Maya leaned forward. "And that is why everything we have been talking about tonight matters so much. The investors you will be sitting across from are going to want to hear you speak fluently about EBITDA margins, revenue multiples, the J-curve, unit economics. They are going to stress-test your assumptions. If you cannot walk them through the value creation bridge we just drew on this napkin, they will move on to the next pitch."

Sarah nodded slowly. "So this is not just academic. I need to internalize all of this before I walk into a single investor meeting."

"You need to own it," Maya said. "Seed investors will give you some latitude—they are betting on the team and the thesis. But Series A investors

will expect real numbers and a credible path to the margins we have been discussing. The founders who raise successfully are the ones who understand the economics as well as the bankers on the other side of the table. And by successfully, I mean the right amount of money under the right conditions. There are a lot of zombie companies out there that took the wrong amount at the wrong time under the wrong terms. That is a trap you do not want to fall into."

2.5.1 *The Politics of Herding Cats*

"Partnerships," Maya said before taking a deep breath.

"Tell me about it," Sarah said.

Maya shook her head. "Even from the banking side, you hear the stories. The economics look compelling on paper. Getting partners to vote for them is another matter."

"Real transformation requires a supermajority vote in many partnerships. That means convincing partners with fundamentally different time horizons, risk tolerances, and economic interests to agree on a multi-year investment that will reduce their near-term income."

"Generational conflict is real. Partners within five years of retirement have limited economic incentive to fund a transformation they will not benefit from. They prefer maximizing distributions today. Younger partners may see the strategic imperative but lack voting power. The middle cohort—partners with 10 to 20 years remaining—are the swing votes."

"Practice group economics differ. Partners billing $1,500 per hour in niche practice areas may see less urgency than other lawyers already losing work to automation. A firm-wide transformation vote pits practice groups with different AI exposure against each other."

"Compensation restructuring is required. Traditional compensation rewards origination and hours billed. AI-native models reward supervision quality, client relationships, and platform contribution. Changing the compensation system while changing the operating model compounds political complexity."

Sarah thought about her old firm. The managing partner who had dismissed AI as another technology wave. The senior partners who blocked any change that might reduce their draws. "That is exactly why I left."

"And that is exactly your advantage." Maya smiled. "You can move while they argue."

"Let me give you an example," Maya said. "In a 20-partner firm, you need 13 to 15 votes for a supermajority. If 5 partners are near retirement and 3 run practices less affected by AI, you start with 8 potential opponents. The remaining 12 must nearly unanimously support transformation."

"That is not a management decision," Sarah said. "That is a political campaign."

"Exactly. I advised a firm recently that had perfect economics for transformation. The managing partner built a compelling case. The investment committee approved the plan." Maya shook her head. "The partner vote failed. They are already losing associates to other firms, and their revenue per lawyer dropped by year's end."

"What happened to the managing partner?"

"She left to join one of the competitors."

Maya leaned forward. "Here is what that story means for you. Every month those firms spend arguing about transformation is a month you spend compounding data and winning their clients. But that window is not permanent. Eventually one of them will get the vote right, or a well-capitalized competitor will enter your segment. You have maybe two to three years where the incumbents are still paralyzed. After that, the race gets harder."

Sarah felt the weight of that. Two to three years. It was both a long time and no time at all.

2.6 Moats, Markets, and Competition

The conversation had turned heavy. Maya ordered dessert to lighten the mood.

"Let me tell you about something that makes me optimistic about firms like yours," she said. "The data moat."

"Yeah, that moat is what we want to build," Sarah said. "I am working on describing that to investors in a way that is not too technical but conveys the key ideas."

Maya laughed. "Fair. An AI-native firm accumulates data from every engagement. Every document processed, every decision made, every client preference captured, every edge case handled."

"I did get a computer science minor when we were in undergrad," Sarah said. "Obviously the cutting-edge models today are way better than what I studied back then, but we are going to develop a whole layer cake of AI techniques to increasingly deliver high-fidelity outputs for various tasks. The most important idea is that whatever we are doing is going to compound."

"Exactly," said Maya. "You need to construct an organization where each week, month, and year of operation adds more data. Competitors cannot easily replicate data accumulated over years. Even if they license the same underlying AI models, they lack all of that very specific information that will drive the highest-quality performance."

Sarah thought about her own firm's data. Four months of contracts, memos, and client communications. Not much yet. But growing every day.

"How do you value a data moat?"

"That is the hard question." Maya stirred her coffee. "Traditional valuation does not capture it well. But when I am advising buyers, I tell them to look for firms where the AI performance improves measurably over time. That is the signal that the data advantage is real."

Sarah had one more question. "What happens when everyone transforms?"

Maya nodded slowly. "This is the question no one in my world has a good answer for yet."

"If AI enables 3 times productivity gains and every firm adopts it, three scenarios are possible."

"Winner-take-most. First movers capture market share before competitors catch up. Data advantages compound—the firm with 50,000 processed contracts has training data that a new entrant cannot replicate in less than years. Network effects emerge as the platform learns client preferences, regulatory patterns, and practice-specific nuances. Clients become sticky because switching means losing the institutional knowledge the platform has accumulated about their business. A few AI-native firms dominate their segments while traditional firms bleed clients. In this scenario, early transformers capture extraordinary value—but the window to establish a position closes fast."

Sarah's pulse quickened. "That is the scenario I am betting on. At least in my corner of the market."

"Maybe. But there is a second possibility."

"Market concentration. Several AI-native firms emerge, each serving different market segments or practice areas. One firm dominates mid-market contract review. Another owns healthcare regulatory compliance. A third captures commercial due diligence. Margins are higher than traditional but not as high as early projections because competition within each segment erodes pricing power. The winners differentiate on domain expertise, client relationships, and data quality rather than raw AI capability—since the underlying models are increasingly commoditized. And there is a natural ceiling on concentration that does not exist in software: conflicts. The more clients you take on in a given sector, the more conflicts you generate. A firm that represents three pharmaceutical companies in regulatory matters will eventually be conflicted out of representing a fourth. That structural constraint means legal markets resist the kind of winner-take-all dynamics you see in technology. Value is captured by efficient operators who combine strong technology with genuine professional judgment. This is probably the most likely outcome in most practice areas."

"The conflicts point is interesting," Sarah said. "I had not thought about that as a structural cap on concentration. In software you want every customer. In law, every new client is a potential conflict with an existing one."

"Exactly. Which is why I said this is probably the most likely outcome. But there is a third scenario, and it is the one you should lose sleep over."

"Commodity race to the bottom. AI democratizes quality to the point where every firm offers similar AI-powered services. The technology becomes table stakes—like having a website or using email. Competition shifts almost entirely to price. Margins compress toward traditional levels or worse, because firms now carry technology costs on top of reduced but still significant labor costs. Value flows to clients in the form of lower fees, and technology vendors—the model providers, the infrastructure companies—capture most of the economic benefit. This is the nightmare scenario for anyone investing in AI-native professional services."

"Which scenario do you think will happen?"

Maya paused. "The honest answer is that it depends on where you sit in the market. And that is actually the good news for you."

She set down her fork. "The legal market is not one market. It is hundreds of markets layered on top of each other. The Am Law 100 is one world—Skadden, Kirkland, Wachtell—firms doing bet-the-company M&A, securities litigation, and cross-border regulatory work where a single matter can generate tens of millions in fees. You are not competing with them. You do not need to."

"Because the rest of the market is enormous," Sarah said.

"Enormous and deeply fragmented. There are tens of thousands of law firms in the United States. The vast majority are small—fewer than 10 lawyers. Mid-market firms like your old one serve regional clients, handle commercial contracts, regulatory compliance, employment disputes, real estate transactions. The work is important but it is not bet-the-company. It is the kind of work where clients are already price-sensitive, already frustrated with the billable hour, and already looking for alternatives."

"That is exactly the market I am targeting."

"And it is a market where AI-native economics are most disruptive. The elite firms can absorb AI as a cost of doing business—their clients pay for judgment and relationships, not efficiency. But in the mid-market and below, clients are choosing between your AI-enabled fixed-fee work and someone else's open-ended billable-hour engagement. The value proposition is obvious."

Maya continued. "The fragmentation also means there is no single incumbent to unseat. You do not need to beat Skadden. You need to win enough mid-market clients in your practice areas to build the data moat and compound from there. A firm doing $5 million in revenue in the legal market is invisible at the industry level—but it can be highly profitable and growing fast."

"So the winner-take-most scenario is unlikely across the whole market."

"Across the whole market, yes. But within specific segments—mid-market contract review, regulatory compliance for healthcare companies, commercial due diligence—a well-positioned AI-native firm can absolutely become the dominant player. The question is not whether you can win the legal market. It is whether you can win your corner of it and then expand from there."

Sarah thought about Candor's first clients. The mid-market tech companies. The healthcare firms navigating regulatory complexity. None of them had ever considered hiring Skadden. They needed competent, efficient, affordable legal work—and they needed it faster than traditional firms could deliver.

"Is that why you agreed to have dinner with me?"

Maya was quiet for a moment. "Partly. I also—" She paused, and for the first time all evening the confident banker seemed to search for words. "I spend my days telling other people what their companies are worth. I model their futures, stress-test their assumptions, negotiate their exits. And sometimes I wonder what it would feel like to actually build something instead of just putting a price on it." She smiled, but it did not quite reach her eyes. "So yes, I agreed to have dinner with you because you are my friend and I miss the old days. But also because watching someone do the thing I advise on—there is a part of me that is a little envious."

Sarah had not expected that. "You could still build something."

"Maybe. But I am good at this, and the comp is hard to walk away from." Maya laughed, and the banker was back. "Which, ironically, is exactly the problem we were just discussing about partnership economics."

2.7 Anatomy of a Margin Machine

Dessert arrived. Maya pulled out a final diagram.

"Hey, let me leave you with the heuristic framework I actually use when evaluating these firms. Everything we have discussed tonight connects to this."

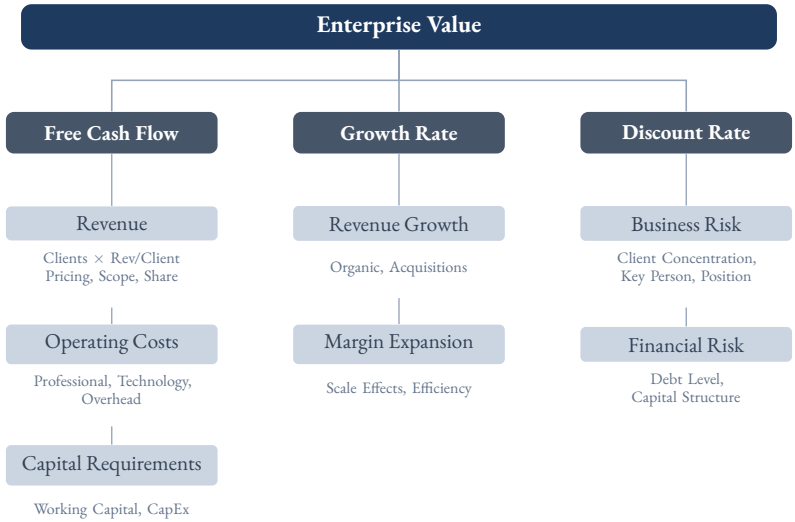

Figure 2.3: Value Driver Tree: How operational decisions flow to enterprise value

"Every decision you make—strategy, services, organization, technology, clients, talent—flows through this tree to affect enterprise value."

Sarah studied the framework. "Revenue growth comes from AI enabling me to serve more clients, win more deals, expand existing relationships."

"Right. Cost reduction operates through professional costs declining via productivity gains and overhead shrinking through automation."

"Productivity improvement is the single largest lever," Sarah said, tracing the logic. "If AI makes my team 3 times more productive, I can triple revenue with the same headcount or maintain revenue with one-third the cost."

Maya held up a hand. "Careful. That math has a hidden assumption. How does your old firm price its work?"

"By the hour. Almost everything."

"So what happens when AI makes a contract review three times faster?"

Sarah saw it immediately. "You bill one-third the hours. Revenue drops."

"Exactly. The productivity paradox. Under hourly billing, getting faster actually shrinks your top line. Every efficiency gain flows straight to the client." Maya tapped the revenue node on the diagram. "The entire margin thesis we just discussed depends on decoupling price from hours worked."

"So you have to move much of the work to some other economic model. Fixed fees, creating a technology surcharge or access license, or you must significantly increase hourly rates but deliver massive efficiency gains per unit of time."

"Listen, it is hard in professional services to fully move away from hourly pricing. Clients are used to it, procurement departments are built around it, and for genuinely unpredictable work it can be the fairest model for both sides. But at a minimum, where you have repeatable, well-scoped work—and AI-native firms will have a lot of it—you should be pricing on value delivered, not hours consumed. Fixed fees, subscriptions, outcome-based arrangements. That is where the margin expansion actually lives. If you are still selling hours, AI just makes you a cheaper law firm. If you are selling outcomes, AI makes you a more profitable one."

Sarah nodded. "So the pricing model is not a downstream decision. It is a prerequisite."

"It is the prerequisite. Get that wrong and everything else in this tree collapses."

Sarah traced the logic back through the tree. "So what happens to margins as productivity scales? If my team moves from 1x to 2x to 3x productivity?"

"The relationship is not linear—it is convex. The first doubling of productivity is transformative. A firm starting at 10 percent EBITDA margin can jump to 25 or even 40 percent, depending on how aggressively it reduces headcount versus grows revenue. But each additional increment yields diminishing returns because you are squeezing cost out of a base that is already lean. The practical ceiling is somewhere around 50 percent EBITDA, even at 4x productivity. And

most firms will not pursue a pure cost-reduction path—they will blend efficiency gains with revenue growth, which means margins rise more gradually but the top line grows alongside them."

Maya picked up her pen one last time. "And it is not just fewer junior lawyers. It is fewer lawyers at every level to deliver the same large-scale engagement. A lawyer with a reasonable amount of experience and the cognitive dexterity to refactor herself to work alongside AI can potentially match the output of someone with significantly more years of practice. It is not a guarantee, and it is going to vary from task to task. But we are already seeing this in the pitches from AI-native firms across professional services—accounting, procurement, management consulting. The ones that figure it out are fielding smaller, flatter teams that punch well above their weight."

Sarah took a photo of the napkin with her phone. "I am keeping this."

The bill arrived. Maya grabbed it before Sarah could move.

"Expense account," she said. "Consider it an investment in deal flow."

Sarah laughed. "You think I am going to be a deal?"

"I think you are either going to be a deal or a competitor to my clients. Either way, I want to stay close." Maya signed the receipt. "I also think you should talk to some investors I know. The economics make sense. You need capital to scale."

"Are you offering to make introductions?"

"I am offering to make introductions to people who will be skeptical but fair. They will push you harder than I did tonight. If you can hold your own and handle the heat, they will write checks."

Sarah sat with that for a moment. It was not that many months ago she had been a senior associate at a traditional firm, watching a managing partner present Microsoft Copilot as 'go to the moon' style innovation. Now she was sitting in a Manhattan restaurant, being offered introductions to investors by an investment banker who thought her firm could eventually be worth hundreds of millions of dollars.

"Thank you," she said. "For all of this."

Maya smiled. "How do you feel?"

Sarah thought about it. "More confident. And more scared."

"Good." Maya gathered her things. "That is exactly the right way to feel."

They walked out into the Manhattan night. The city hummed around them—taxis, pedestrians, the ambient energy of ambition and capital and risk.

"One more thing," Maya said as they waited for Sarah's car. "The frameworks I showed you tonight—they are useful. But they are also wrong. All models are wrong. Some are useful—George Box."

"I remember that from our econ class all those years ago," said Sarah.

Maya hugged her and sent her off with one last thought. "Go build something. Then come back and tell me I underestimated the margins."

Sarah watched her friend disappear into the crowd. She had come to New York with questions. She was leaving with napkin sketches, wine-stained placemats, and something that felt like a plan. All models are wrong, Maya had reminded her. But the model in Sarah's head—the one she had been assembling all evening from Maya's frameworks and her own conviction—felt less wrong than anything she had carried before.

What came next would be harder. It was time to turn the napkin sketches into a real company.

Coming Next Week

Chapter 3
Uncertainty, Structure, and Decision

How should firms structure themselves to navigate transformation under genuine uncertainty? Chapter 3 introduces the Jurisdictional-Structure Matrix and the Scenario Matrix—frameworks for making structural and strategic decisions when the future is unknowable. Sarah confronts the question every founder dreads: which jurisdiction, which entity structure, and which regulatory path gives her firm the best chance of surviving long enough to prove the thesis.

Subscribe at `https://theainativefirm.com` to receive new chapters as they release.

About the Authors

Daniel Martin Katz, PhD, JD is Professor at Illinois Tech–Chicago Kent College of Law and Academic Director of the Bucerius Center for Legal Technology & Data Science. Named by the *Financial Times* as one of the top 20 legal market shapers of the past twenty years, his research focuses on legal analytics, legal technology, and the future of the legal profession.

Michael J. Bommarito II, MSE, MA is a serial entrepreneur, researcher, and adjunct professor with twenty-five years of industry experience. He has founded and led multiple companies at the intersection of artificial intelligence and professional services.

Jillian Bommarito, CPA, CIPP/US/E is an advisor, risk and governance expert, and one of the first certified AI auditors in the world. She brings deep expertise in compliance, privacy, and the governance challenges of deploying AI in regulated industries.

Get the Full Book

`https://theainativefirm.com`