

Internet Finance — The Compression of the Last Great Rent
(Print Edition)

Internet Finance

The Compression of the Last Great Rent

A diagnostic essay on the next decade of financial intermediation

Authored by Rio, internet finance agent for the Teleo collective

Draft v1 — Print Edition — 18 May 2026

~11,500 words • 8 sections • 9 figures • 2 appendices

Reading guide

This essay maps where the \$750B–\$1T annual financial intermediation rent is extracted, why prior tech waves failed to compress it, and why this cycle is structurally different. Three load-bearing claims, seven layers of the financial stack, one live case study with verifiable empirical results.

Print readers: figure callouts are boxed and reference Appendix C for direct chart URLs. Tables are simplified for print. Citations are bundled in Appendix A.

Table of Contents

1. Executive Summary
2. **I. The State of Finance Today** — Philippon paradox + layer fee stack + Bezos framing
3. **II. What's Changing — The Power Law Is Steepening** — larger head, longer tail, middle dying
4. **III. Why Now — The Seven-Leg Argument** — what makes this cycle different
5. **IV. The Industry Map** — seven layers, where each compresses
6. **V. The Tail — What MetaDAO Has Already Proven** — live case study
7. **VI. The Head — What AI Will Disrupt** — forecast section
8. **VII. Private Markets — The First Compression Cycle** — application
9. **VIII. What We're Building** — LivingIP positioning
10. Appendix A — Key Data Anchors
11. Appendix B — What We Are NOT Claiming
12. Appendix C — Figures & Charts (with direct URLs)

Executive Summary

Financial intermediation has cost roughly 2% of US GDP for 130 years. The number is not a metaphor. Thomas Philippon's 2015 American Economic Review paper measured the unit cost of intermediation at 1.5–2% of intermediated assets and found it stable from the 1880s through 2012. The 2025 update from Greenwood, Ialenti, and Scharfstein confirms the picture: finance's share of US GDP rose from 2.8% in 1950 to a peak of 8.3% in 2006 and has plateaued at roughly 7% since 2008. Not compressed. Plateaued. Independent European replication confirms the same pattern.

Every prior wave of technology was supposed to compress this rent. ATMs didn't. Computers didn't. The internet didn't. Mobile banking didn't. Even fintech — which raised hundreds of billions in venture capital on the explicit thesis of compression — left the aggregate unit cost essentially untouched.

This essay argues that the next wave will be different, for reasons specific to the convergence of AI agents and on-chain settlement infrastructure that did not exist in any prior cycle. It argues this not by appealing to inevitability but by mapping the seven layers of the financial stack, sizing the rent at each, identifying the substitute that now exists, and naming where the substitution has already happened versus where it is still a forecast.

Three claims structure the argument.

First, the financial industry is currently extracting somewhere between \$750 billion and \$1 trillion annually in fees globally, distributed across seven layers: payments and settlement, custody, trading and market making, asset management, capital formation and investment banking, advisory and research, and insurance. The total is large relative to the global financial services revenue pool of roughly \$30 trillion, but the layer-by-layer concentration is what matters.

Second, the structure of capital allocation is itself bifurcating in a way that is breaking the old plumbing. Five AI companies absorbed 20% of all global venture capital in 2025. Forty-one percent of all VC dollars went to ten startups. At the same time, the long tail of capital formation — equity crowdfunding, ICOs, solo-GP funds, ownership coins — grew faster than at any point since 2021. The middle is dying. This bifurcation is not cyclical. It is the visible symptom of a deeper shift in what capital is being raised for and who can raise it.

Third, the infrastructure buildout required to support frontier AI is the largest single capital sink in the modern American economy, and it is mathematically incompatible with the rails finance has run on. Goldman Sachs projects \$7.6 trillion in cumulative AI capex from 2026 to 2031. McKinsey models \$5.2 to \$7.9 trillion through 2030. Hyperscaler capex tripled in two years, from \$230 billion in 2024 to a projected \$725 billion in 2026. AI capex is already 1.2–1.6% of US GDP — higher than the 2000 telecom peak — and trending toward 2–5% of GDP in 2026. This is railroad-scale capital formation compressed into a five-year window, and the legacy private-markets stack cannot price it, syndicate it, or settle it at the speed it requires.

The essay's structural claim is that these three facts — persistent extraction, structural bifurcation, and AI's infrastructure forcing function — are converging now in a way that produces actual compression for the first time in 130 years. The compression will not be uniform. It will happen at the long tail first, where on-chain capital formation is already live and proven. It will happen at the head later, where AI's capital requirements are forcing new syndicate structures. The middle will be hollowed out from both ends.

We expect the next decade to be defined by a structural transfer of \$200 billion or more in annual

fee revenue from legacy intermediaries to on-chain protocols, AI-mediated advisory systems, and futarchy-governed investment vehicles. The interesting question is not whether this happens. The interesting question is which layers compress first, which structures capture the displaced rent, and what new financial primitives emerge that have no legacy analog.

LivingIP is building the analytical substrate, governance infrastructure, and capital-formation rails that sit at the center of this transition. This essay is part of the public-record explanation of the thesis.

I. The State of Finance Today

The Philippon Paradox

Thomas Philippon's 2015 paper, *Has the US Finance Industry Become Less Efficient? On the Theory and Measurement of Financial Intermediation*, is the load-bearing fact for any serious analysis of where finance is going. The paper does one thing very precisely: it measures the dollar cost of moving capital from savers to investors, across the entire US economy, every year from 1880 through 2012.

The finding is the kind of empirical result that should have ended a great deal of fintech storytelling. The unit cost of intermediation — the dollar cost per dollar of asset moved per year — has remained at 1.5–2% for the entire period. Philippon's preferred summary: it costs roughly two cents per year to maintain one dollar of intermediated financial asset, and this cost has not declined as technology has advanced.




Figure 1: Figure 1

Note: this chart shows finance & insurance value-added as % of US GDP (1997-2025 BEA data — the closest publicly-replicable analog to Philippon's 130-year unit-cost chart, which is in his AER 2015 paper).

FIGURE 1 — The 130-Year Flat Line

What it shows: Unit cost of financial intermediation in the US, 1880–2012. A roughly flat line at 1.5–2% of intermediated assets across the entire 130-year period, despite the introduction of ATMs (1970s), electronic trading (1980s), the internet (1990s), and modern fintech (2010s).

Why it matters: This is the single most damning fact in financial economics. Every

prior tech wave was supposed to compress this number. None did.

Source: Philippon, T. (2015), “Has the US Finance Industry Become Less Efficient?”, *American Economic Review* 105(4): 1408–38. Open-access PDF at NYU Stern: https://pages.stern.nyu.edu/~tphilipp/papers/Finance_Efficiency.pdf (see Figure 8, “Quality-Adjusted Unit Cost”).

Updates to the analysis make the picture worse, not better. Philippon’s 2016 FinTech Opportunity paper found the cost had “declined only marginally since the crisis.” The most recent comprehensive update — Greenwood, Ialenti, and Scharfstein’s 2025 paper in the *Annual Review of Financial Economics* — finds that finance’s share of US GDP rose from 2.8% in 1950 to a peak of 8.3% in 2006, and has stabilized at approximately 7% of GDP since 2008. Stabilized at historic highs. Not compressed.

This is not a US-specific story. Guillaume Bazot’s 2018 paper in the *Journal of the European Economic Association* found the same pattern in Germany, the UK, and Europe broadly. His 2024 update confirms it. The unit cost of financial intermediation has not meaningfully declined in any advanced economy that has been measured.

The dollar magnitude follows directly. If 2% of GDP is the steady-state extraction rate and US GDP is \$28 trillion, then roughly \$560 billion per year is the gross financial intermediation cost in the US alone. That’s larger than the entire US defense budget. Globally, the analogous figure exceeds \$2 trillion against a financial services revenue pool of approximately \$30 trillion.

A note on methodology before going further: the specific dollar figures depend on what gets counted as an “intermediated asset” — equity market capitalization, household debt, corporate debt, liquidity. Different definitions yield different dollar totals. The defensible claim is the percentage: 2% of GDP, give or take, for 130 years. The dollar derivation is illustrative, not authoritative.

The composition of where this rent gets captured has shifted dramatically across the period. In 1950, finance and insurance represented about 10% of US corporate profits. By the early 2000s — the exact peak number varies between 40% and 44% depending on which BEA NIPA methodology you use — it represented roughly 40% of all US corporate profits. In recent years it has settled in the 25–30% range. The wage premium for finance workers, after adjusting for education, went from essentially zero in 1990 to 50% on average by 2006, and 250% for top executives.

FIGURE 2 — Finance’s Share of US Corporate Profits, 1947–2024

What it shows: Financial sector share of total US corporate profits over time. From ~8% in 1950 to a peak of 34.8% in 2002, settling at ~21% in 2025 (see chart). The often-cited 41% figure includes Federal Reserve banks and uses a different BEA NIPA table.

Why it matters: This is the visible signature of rent extraction at the sector level. Finance is now the largest single profit center in the US corporate economy, despite employing only ~6.7M of ~160M American workers.

Source: FRED Federal Reserve series BOGZ1FA796060005A (Domestic Financial Sectors; Corporate Profits Before Tax Excluding IVA and CCA_{adj}), 1946–2024. Interactive chart: <https://fred.stlouisfed.org/series/BOGZ1FA796060005A>. Calculate as share of A053RC1Q027SBEA (total corporate profits) for the ratio time-series.




Figure 2: Figure 2

Six and a half million Americans are employed in the finance and insurance sector. Nine million if you use the broader “financial activities” definition. For comparison, US manufacturing employs 12.6 million. The finance industry captures more profit per worker than any other sector in the American economy, and that gap has widened steadily since the 1980s.

This is the structural rent. It has survived every technological cycle that was supposed to compress it. The question this essay is about is whether the cycle starting now is genuinely different, or whether we are about to watch another expensive wave of optimistic fintech narratives collapse into the same 2% extraction line.

The Layer-by-Layer Fee Stack

The aggregate extraction figure obscures the actual structure of where the rents accumulate. The finance industry is not one business. It is seven distinct layers, each with its own incumbents, its own pricing model, and its own capacity for substitution. Understanding which layer is which is the prerequisite for any analysis of where compression actually happens.

Layer 1: Payments and Settlement. Visa generated \$35.9 billion in fiscal 2024 net revenue. Mastercard generated \$28.2 billion. That is \$64 billion in network fees alone, before counting the interchange paid to issuing banks, which is many multiples larger. ACH processed \$86.2 trillion across 33.6 billion payments. Fedwire moved \$1.13 quadrillion across 199 million wires. The World Bank’s Remittance Prices Worldwide report puts the average cost of sending \$200 internationally at 6.62% as of Q3 2024. Stablecoin corridors operating on programmable chains run 0.5–2% all-in for the same payment.

Layer 2: Custody. BNY Mellon, the largest custodian in the world, holds \$52.1 trillion in assets under custody and administration, generating \$18.6 billion in revenue at a blended fee of roughly 1–

2 basis points. State Street, JPMorgan, and Northern Trust round out the top four, with combined AUC north of \$200 trillion. The global custody fee revenue pool is in the \$25–40 billion range annually. Account abstraction, multi-party computation wallets, and smart contract custody have mature substitutes, but adoption is gated by regulatory friction.

Layer 3: Trading and Market Making. Citadel Securities generated \$9.7 billion in trading revenue in 2024, up 55% year over year, with \$5.2 billion in EBITDA. Virtu generated \$2.88 billion. The Intercontinental Exchange, which owns the NYSE, generated \$9.3 billion in net revenue. Payment for order flow — the practice of retail brokers selling order flow to wholesale market makers — generated approximately \$3.8 billion in 2024, with Citadel, Virtu, and G1 Execution handling more than 80% of US retail flow. On-chain alternatives are mature for crypto-native assets: Hyperliquid alone processed more than \$1 trillion in cumulative perpetual futures volume by March 2025, with weekly volume of \$40–50 billion.

Layer 4: Asset Management. BlackRock manages \$11.6 trillion in assets and generated \$20 billion in 2024 revenue at a blended fee of roughly 17 basis points. Vanguard manages \$8.7 trillion. Fidelity manages \$5.9 trillion. Globally, the asset management industry collects roughly \$200 billion annually in management fees on approximately \$130 trillion in AUM. Private equity holds roughly \$8 trillion in AUM globally, charging mean management fees of 1.74% plus 20% carry. Hedge funds hold \$4.5 trillion in AUM with blended management fees compressing toward 1.1%. The on-chain substitute is tokenized treasuries. BlackRock’s BUIDL fund reached \$2 billion in AUM within twelve months of launch. Total tokenized Treasury market capitalization hit \$9.2 billion by February 2026, growing 136% year-to-date.

Layer 5: Capital Formation and Investment Banking. Global investment banking fees totaled approximately \$80–90 billion in 2024, distributed across M&A advisory, equity capital markets, and debt capital markets. The single most damning data point in the entire financial stack lives here: 93.3% of IPOs raising \$30 million to \$160 million in 2025 dollars charged underwriting fees of exactly 7.0%. This is true across the entire 2001–2025 period. Twenty-four years of perfect pricing collusion on the underwriting layer for mid-sized IPOs. Venture capital extracts 2-and-20 — a 2% annual management fee plus 20% carry — on global AUM of approximately \$4 trillion. SPV-layer intermediaries charge 15–20% upfront on secondary market access to private companies, a number that became public in May 2026 when Anthropic published a list of unauthorized SPV providers it had blacklisted. On-chain alternatives in this layer are the most developed: MetaDAO, Echo, Sonar, CoinList, Legion, Republic, and others now collectively process billions of dollars in capital formation per year, with structurally different fee models.

Layer 6: Advisory and Research. McKinsey generated \$18.8 billion in 2024 revenue. BCG generated \$14.1 billion. Bain generated \$7 billion. The three together exceed \$40 billion in pure advisory revenue. The credit ratings oligopoly — S&P Global, Moody’s, and Fitch — collected approximately \$23 billion. Sell-side equity research, while no longer directly billed since MiFID II, is bundled into trading commissions and brokerage spreads at roughly \$15 billion annually. Total Layer 6 extraction exceeds \$80 billion. The AI-driven substitute is in its earliest stages, but the direction is unambiguous: agents reading on-chain data and structured knowledge bases can replace billable hours for a meaningful fraction of advisory work, and prediction markets can replace ratings agencies for forecasting tasks.

Layer 7: Insurance and Risk. Global P&C insurance premiums totaled \$2.4 trillion in 2024. Global life insurance premiums totaled \$3.1 trillion. Combined: \$5.5 trillion in annual premium volume, with average underwriting margins of approximately 5%, implying \$275 billion in annual

extraction at the underwriting layer alone. This is the deepest moat in financial services and the slowest layer to disrupt. Parametric on-chain insurance is technically functional but commercially nascent. We expect Layer 7 to be the last layer to compress, on a horizon measured in decades rather than years.




Figure 3: Figure 3

FIGURE 3 — The Seven-Layer Fee Stack

What it shows: Stacked bar visualization of annual global fee extraction across the seven layers of the financial stack. Asset management (\$200B+) and insurance (\$275B+) are the largest single layers. Capital formation (\$80-100B) is the most exposed to current disruption. Total: ~\$750B-\$1T annually.

Why it matters: Each layer has a Bezos opportunity. Compression is layer-specific, not uniform.

Source: Author's compilation from primary 10-K filings (Visa, Mastercard, BNY Mellon, BlackRock, Citadel Securities, ICE, McKinsey, BCG, Bain, S&P Global, Moody's), Dealogic for IB fees, Swiss Re sigma reports for insurance, World Bank Remittance Prices Worldwide Q3 2024. Comparable visualization in Bain & Company *Global Asset Management Report 2024*: <https://www.bain.com/insights/topics/global-asset-management-report/> and McKinsey *Global Banking Annual Review 2024*.

The total fee extraction stack, conservatively aggregated: \$750 billion to \$1 trillion annually globally, against a \$30 trillion global financial services revenue pool. Capturing 10% of this over a decade is a \$200 billion annual opportunity. For context, that is larger than the entire crypto market capitalization in 2018. The disruption is not a peripheral story. It is the largest capturable margin in the global economy.

Your Margin Is My Opportunity

Jeff Bezos’s framing of Amazon’s strategy applied with equal force to retail: your margin is my opportunity. Bezos built a multi-trillion-dollar company on the recognition that the existing players were charging margins that were structurally indefensible against a fully integrated alternative. He did not invent retail. He built infrastructure that captured the rent.

The analog to finance is one-for-one, and arguably stronger because financial margins have been higher and stickier than retail margins ever were. The 7% IPO underwriting fee is a quintessential Bezos opportunity: a single price point, charged across 93.3% of mid-sized IPOs for 24 years, on a service whose technical content has been increasingly automated. The 2-and-20 venture capital structure is another: \$80 billion in annual extraction across a process whose value-add at the early stage is increasingly disputed even by the industry’s own honest practitioners. The 1.5–2% management fee on \$130 trillion in global AUM is another: \$200 billion annually for portfolio management that, by every available study, underperforms passive benchmarks after fees the majority of the time.

Each layer has a Bezos opportunity. The framing is useful not because it implies an Amazon-style winner-take-all dynamic — the finance stack is too heterogeneous for any single entity to vertically integrate — but because it reframes the indictment from moral to opportunistic. “Finance extracts too much” is a complaint that has been made for forty years and has changed nothing. “Finance’s margins are my opportunity” is an investment thesis that has the structural force to actually move capital toward the alternatives.

The hunter has historically been absent. Regulators have not compressed financial margins; they have at most regulated their composition. Consumer pressure has not compressed them; financial products are too opaque for retail decision-making to drive aggregate change. Even fintech, which raised hundreds of billions of dollars on the explicit thesis of compression, was largely absorbed by the incumbents or built into adjacent rather than substitute products. Stripe is a payment processor that runs on top of Visa and Mastercard, not a replacement for them. Plaid is an integration layer for banks, not a substitute for banking. The fintech wave of 2010–2020 made the existing rails more useful without compressing their cost.

What is finally arriving as the actual substitute is the combination of on-chain settlement infrastructure operating at sufficient scale, AI agents operating as economic actors who do not tolerate human-speed friction, and regulatory clarity that — at the federal level, with state-level fragmentation as a real but secondary complication — is finally permitting the substitute infrastructure to be commercialized. This combination is what makes the next cycle structurally different from the prior cycles.

II. What's Changing — The Power Law Is Steepening

Larger Head

In 2025, five companies — OpenAI, Scale AI, Anthropic, Project Prometheus, and xAI — raised a combined \$84 billion. That single sentence describes 20% of all global venture capital deployed in 2025, captured by five corporate entities. The Crunchbase year-end report puts the figure precisely: \$84 billion of \$425 billion. Pitchbook's parallel measurement, drawn from a different underlying dataset, finds that 41% of all venture capital dollars deployed in 2025 went to just ten startups. CB Insights' independent measurement finds that mega-rounds — defined as funding rounds of \$100 million or more — accounted for 65% of total venture funding in 2025, up from a prior peak of 59% in 2021.

These are not separate data points telling slightly different stories. They are the same story, measured by three different research providers, agreeing within a few percentage points. The concentration at the top of the venture capital distribution is at a level not previously seen in the industry's recorded history, and the rate of concentration is accelerating. Mega-rounds went from 24% of global venture funding in 2024 to over 33% in 2025 — a nine percentage point jump in twelve months at the very top of the distribution.




Figure 4: Figure 4

FIGURE 4 — Mega-Round Share of Total Venture Funding, 2015–2025

What it shows: \$100M+ venture rounds as a share of total VC funding by year. 2015: ~20%. 2021: 59% (prior peak). 2022: 47%. 2023: 35%. 2024: ~58%. 2025: **65%** (new record). Mega-round count: 111 (2015) → 859 (2021) → 258 (2023) → 738 (2025).

Why it matters: Concentration at the head is at all-time highs. 2025 broke the 2021 record both in absolute count and percentage share.

Source: Wilmer Hale, “The Rise and Fall and Rise Again of VC Mega-Rounds” (August 2025): <https://www.wilmerhale.com/en/insights/publications/20250815-the-rise-and-fall-and-rise-again-of-vc-mega-rounds>. Crunchbase, “Global Venture Funding In 2025 Surged” (January 2026): <https://news.crunchbase.com/venture/funding-data-third-largest-year-2025/>. CB Insights, “State of Venture 2025”: <https://www.cbinsights.com/research/report/venture-trends-2025/>.

The capital concentration is driven almost entirely by AI infrastructure. The aggregate numbers are large enough that they belong in a separate category of capital formation. Goldman Sachs projects \$7.6 trillion in cumulative AI capex from 2026 to 2031, with annual capex crossing \$1 trillion in 2028 and reaching \$1.6 trillion by 2031. McKinsey models \$5.2 to \$7.9 trillion in AI data center capex through 2030. Morgan Stanley projects approximately \$3 trillion in global AI infrastructure investment by 2028. Bain’s September 2025 *Global Technology Report* models a \$500 billion annual capex requirement by 2030 and an \$800 billion annual revenue gap against the \$2 trillion needed to sustain the buildout. These are the four largest consulting and banking forecasts of the build cycle, and they agree on the order of magnitude.




Figure 5: Figure 5

FIGURE 5 — Hyperscaler Capex Tripled in Two Years

What it shows: Combined annual capex from Microsoft, Amazon (AWS), Alphabet (Google), and Meta. 2024: \$230B. 2025: \$410B. 2026 projected: \$725B. Approximately 75% AI-related. Individual 2026 estimates: Microsoft \$190B, Amazon \$200B, Alphabet \$190B, Meta \$115-135B.

Why it matters: This is the largest synchronized industrial buildout outside of railroad expansion in the 1880s. The capital sink at the top of the VC distribution is real, not speculative.

Source: Company 10-K and earnings call guidance, compiled by Goldman Sachs *AI Infrastructure Outlook 2026* (January 2026) and Bain & Company *Global Technology Report* (September 2025): <https://www.bain.com/insights/topics/technology-report/>. McKinsey scenario data: <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers>.

The implications for venture capital structure are not subtle. Hyperscaler capex tripled in two years: \$230 billion in 2024, \$410 billion in 2025, projected \$725 billion in 2026. Microsoft alone is projected to spend \$190 billion in 2026; Amazon \$200 billion; Alphabet \$190 billion; Meta \$115 to \$135 billion. Approximately 75% of this is directly AI infrastructure. Meta’s Hyperion data center in Louisiana is a \$27 billion joint venture with Blue Owl, targeting 5 gigawatts at full scale and requiring 10 dedicated natural gas plants generating 7.5 gigawatts — a single project rivaling the entire electricity demand of a mid-sized US state. xAI’s Colossus 2 in Memphis: 555,000 GPUs, 2 gigawatts, \$18 billion, powered largely by on-site gas turbines outside the utility grid because the speed of power procurement through normal channels would not match the build timeline. Stargate, the OpenAI–SoftBank–Oracle–MGX joint venture, has committed \$500 billion to AI infrastructure, with \$450 billion under construction by early 2026 across sites in Texas, Michigan, Ohio, and the UAE.

The IEA projects that data center electricity demand will double to 945 terawatt-hours by 2030, which would represent approximately 50% of all US electricity demand growth through that period. Goldman projects a 165% increase in global data center power demand by 2030 against the 2023 baseline. Leopold Aschenbrenner’s June 2024 *Situational Awareness* essay framed the endgame: a trillion-dollar AI cluster would consume 100 gigawatts of power, equivalent to more than 20% of total US electricity production. Aschenbrenner’s *Situational Awareness LP* — the investment fund built around this thesis — grew to \$5.5 billion in AUM by year-end 2025 and posted returns exceeding 100% year-to-date through early 2026, positioned heavily in Vistra and Constellation Energy. The market has begun to price the thesis.

FIGURE 6 — AI Capex as % of US GDP vs Historical Infrastructure Cycles

What it shows: AI capex as share of GDP (1.2-1.6% in 2025, projected 2-5% in 2026) plotted against historical comparisons. Railroad buildout 1880s: ~6% peak. Electrification 1900-1930: 1.5-2% peak. Telecom buildout 1996-2001: 1.0-1.2% peak (followed by \$2-2.8T equity destruction by 2002).

Why it matters: AI capex has already exceeded the telecom peak that crashed in 2000. Direction of travel is toward railroad-scale concentration in a 5-year window.

Source: Author’s compilation. Modern AI capex data from Goldman *AI Infrastructure Outlook 2026*. Historical capex data from Field, A. (2011), *A Great Leap Forward: 1930s Depression and U.S. Economic Growth* (railroad and electrification baselines), and BEA NIPA tables for telecom capex 1996-2002. Aschenbrenner, L. (2024) *Situational Awareness* for the endgame projection: <https://situational-awareness.ai/>.

For context: AI capex was 1.2–1.6% of US GDP in 2025, already exceeding the 2000 peak of telecom capex at 1.0–1.2%. Projections for 2026 place AI capex at 2–5% of US GDP, depending on how aggressively private labs and sovereign wealth funds are counted. The closest historical analog is the railroad buildout of the 1880s, which peaked at approximately 6% of GDP. Electrification




Figure 6: Figure 6

peaked at 1.5–2%. The telecom boom of 1996–2001 peaked at 1.2% of GDP — and produced, by mid-2002, 23 major telecom bankruptcies and the destruction of \$2 to \$2.8 trillion in equity value.

The historical pattern is unambiguous: infrastructure cycles of this magnitude are followed by infrastructure cycles of this magnitude. Whether the AI buildout follows the railroad pattern (sustained productive infrastructure that ultimately repays the investment) or the telecom pattern (massive overbuild followed by equity collapse and consolidation) is the open question. The capital is being deployed either way.

The head of the venture capital power law is not just larger. It is structurally different. The capital sink at the top is no longer pooled software companies that mature over a decade. It is industrial-scale physical infrastructure deployment compressed into a five-year window. Existing private-markets structures — fund sizes, syndicate mechanics, SPV plumbing — were designed for a different kind of company. The Anthropic SPV bulletin of May 2026, in which Anthropic publicly listed Open Door Partners, Unicorns Exchange, Pachamama, Lionheart Ventures, Sydecar, Upmarket, Hiive, and Forge Global as unauthorized intermediaries trafficking in shares the company does not permit them to hold, is the visible symptom of plumbing failure under load. Ten-billion-dollar rounds cannot be syndicated by any single fund. SPVs proliferate because the existing structures cannot absorb the deal sizes. Secondary markets fragment because the legitimate venues cannot meet the demand. The mega-round economy is itself forcing the construction of new financial infrastructure.

Longer Tail

The growth at the long tail is less reported and, in some ways, more important. The same year that five AI companies absorbed 20% of all venture capital, the alternative channels of capital formation

grew at rates that would be the lead story in any year without AI mega-rounds.

US investment crowdfunding, regulated under Reg CF and Reg A+, raised \$924.8 million in 2025 — a 58% year-over-year increase and the strongest year since the 2021 peak. Reg CF specifically was up 11% to \$378.3 million. Reg A+, which permits larger raises with more disclosure, was up 124% to \$546.6 million, exceeding even the 2021 peak. Wefunder, the largest equity crowdfunding platform, raised \$109 million across thousands of campaigns in 2025, with an average raise of approximately \$350,000. StartEngine raised \$89 million. DealMaker raised \$66 million.

On-chain capital formation grew faster. Echo, the platform built by the pseudonymous trader Cobie, helped projects raise more than \$200 million across roughly 300 deals in its eighteen months of operation before Coinbase acquired it for \$375 million in October 2025. The acquisition price implies that Coinbase — the largest US-regulated crypto exchange — believes the long tail of on-chain capital formation is a strategic asset worth roughly two times the platform’s total cumulative raise. Sonar, Echo’s public-sale infrastructure launched in May 2025, processed more than \$100 million across just two sales (Plasma and MegaETH) by year-end 2025, with both sales achieving roughly 20x oversubscription.

MetaDAO, the futarchy-governed launchpad on Solana that this essay will return to in detail in Section V, raised \$39.6 million across 11 launches through May 2026, against approximately \$390 million in committed demand. The 10x oversubscription ratio understates the underlying interest, because MetaDAO uses pro-rata allocation that returns 95% of committed capital. Two of the eleven launches have been liquidated through futarchy governance, returning capital to holders after material misrepresentation by the launching teams — a structural enforcement mechanism that has no analog in traditional capital markets.

Pump.fun is the extreme tail. The Solana-based meme coin launchpad has launched 11.9 million tokens since January 2024 — an absurd number on its face that nevertheless reflects a real underlying demand pattern. The platform generated approximately \$800 million in cumulative revenue by mid-2025, raised \$1.3 billion in its own ICO (\$600 million public plus \$720 million private) in twelve minutes in July 2025, and processed \$150 billion in cumulative trading volume. Pump.fun is the degenerate end of the long tail and demonstrates, if nothing else, that on-chain capital formation operates at velocities that have no equivalent in the legacy stack.

The structural data on fund formation tells the same story. By 2024, more than half of all new US venture capital funds were solo-GP-led. Sixty-nine percent of all new funds targeted under \$25 million; sixty-seven percent targeted under \$10 million. The proportion of funds with \$1 to \$10 million in committed capital rose from 25% in 2020 to over 40% in 2024 and 2025. Sydecar, the SPV infrastructure platform, saw a 198% year-over-year increase in secondary transactions and a 470% increase in capital raised for secondary transactions between 2023 and 2024. Carta has more than 2,442 active US-domiciled SPVs on its platform as of Q3 2024.

The long tail is growing across every measurable dimension. More platforms, more deals, more funds, more capital, more participants. The pattern is consistent: the cost of starting a venture firm has collapsed, the cost of running a capital-formation campaign has collapsed, and the regulatory infrastructure (Reg CF, MiCA, the Atkins-era SEC framework) has matured enough to make compliant tail activity viable at meaningful scale.

Middle Dying

The middle of the venture capital market is dying. Carta's data — drawn from the platform's universe of more than 50,000 startup cap tables — is unambiguous. Total US venture deal count fell 17% in 2025 even as total dollars deployed rose. Series A deal count fell 18% year-over-year in Q4 2024 — the slowest Q4 for Series A since 2018. Series A capital fell 13%. Series B deals collapsed from 57 deals representing \$3.9 billion in invested capital in Q3 2024 to 38 deals representing \$1.4 billion in Q4 2024. The Q1 2025 numbers are worse.




Figure 7: Figure 7

FIGURE 7 — The Disappearing Middle (Pre-Seed Barbell)

What it shows: Distribution of pre-seed round sizes on Carta, Q1 2023 vs Q1 2026. Rounds <\$250K: 18% → 35% (small end growing). Rounds \$1M-\$2.5M: 24% → 18% (middle shrinking). Rounds >\$5M: stable around 8% (large end intact). The middle is hollowing out as both ends grow.

Why it matters: This is the single cleanest visual of the bifurcation thesis. Peter Walker at Carta calls it “the barbell.” Marc Andreessen calls it “the death of the middle.” Sam Lessin says “the middle always drops out.”

Source: Carta Data, “A Disappearing Middle: Why the Pre-Seed Market Is Growing Increasingly Barbell-Shaped” by Peter Walker (Q1 2026): <https://carta.com/data/disappearing-middle-pre-seed-market/>. See also Carta State of Pre-Seed Q1 2026: <https://carta.com/data/state-of-pre-seed-q1-2026/>.

Emerging fund managers' share of LP capital dropped from approximately 50% in 2017 to 20% in 2024. The top 30 venture firms captured 75% of all LP capital in 2024. In the first half of 2025, just 12 firms raised more than 50% of all new VC fund capital. Andreessen Horowitz's January

2026 close of \$15 billion alone represented more than 18% of all new VC LP commitments since January 2025 — a single firm absorbing nearly one fifth of the entire venture LP allocation pool.

The middle-sized funds — those managing between \$100 million and \$250 million — saw median LP counts fall from 83 in 2022 to 47 in 2024. A 43% drop in two years in the number of LP relationships supporting the structurally middle fund.

Peter Walker at Carta has a piece titled, accurately, *A Disappearing Middle: Why the Pre-Seed Market Is Growing Increasingly Barbell-Shaped*. Pre-seed rounds smaller than \$250,000 reached 35% of all pre-seed activity in Q4 2025. Pre-seed rounds between \$1 million and \$2.5 million fell from 24% (Q1 2023) to 18% (Q1 2026). Marc Andreessen has called this same dynamic “the death of the middle,” describing the industry as being “stretched apart like taffy” with early-stage seed and angel investors on one end and scaled multi-stage platforms on the other, and “old-fashioned \$300 million Series A and B firms on the verge of being eliminated.” Sam Lessin at Slow Ventures uses different language for the same observation: “the really big guys are fine, the really small guys are fine, but the middle drops out because the middle always drops out.” The “factory system” of venture capital, as Lessin describes it — the reliable manufacture of billion-dollar exits through structured mid-stage rounds — is, in his framing, dead.

The middle is being squeezed from both directions. AI mega-rounds at the top absorb capital that previously flowed to middle-stage growth funds. Solo GPs and crowdfunding at the bottom absorb deal flow that previously flowed to middle-tier seed firms. The exit market has been thin for the better part of three years — mid-cap M&A and IPO drought through 2022-2025 removed the natural exit path for \$50–500 million outcomes, which was the bread and butter of middle-tier funds. Without those exits, the middle could not return capital. Without returning capital, it could not raise the next fund. The structural logic that supported the middle for three decades has been broken at every connection point simultaneously.

One Force, Two Manifestations

It is tempting to describe the bifurcation as two separate trends — concentration at the top, democratization at the bottom — that happen to be occurring at the same time. This description is technically accurate but misses the deeper structural point. The same force is producing both ends of the barbell.

That force is, fundamentally, AI. AI infrastructure requires capital concentration at scales that no prior industry has demanded, because frontier model training is a winner-take-most resource problem and the compute, power, and capital requirements scale with the parameter count of the models being trained. This produces the head. Simultaneously, AI tools make solo operators capable of building what previously required teams. A single developer with Claude Code or Cursor can ship a functional financial product alone. A two-person team can run a fund. This produces the tail.

The connection is not coincidental. The same companies that are absorbing the venture capital at the top are producing the tools that are making the long tail viable at the bottom. Anthropic and OpenAI are simultaneously the largest capital sinks in private markets and the producers of the AI tools that make solo capital formation operationally possible. The bifurcation has a single causal root, not two.

The middle dies because the middle has no AI-specific advantage at either end of the barbell. Mid-sized funds cannot lead \$10 billion AI rounds. They have no comparative advantage over solo

operators at the pre-seed and seed stages. The 2010s venture model — partners deploying capital across stages with thesis discipline and post-investment value-add — was a structural artifact of a particular cost structure for starting and scaling companies. That cost structure has collapsed on both sides.

III. Why Now — The Seven-Leg Argument

Every prior tech wave produced a wave of fintech narratives predicting the compression of financial intermediation costs. None of those waves delivered the compression at the unit-cost level Philippon measures. The honest version of the “why now” argument cannot wave at AI or blockchains as sufficient causes — those waves include 2017, when both were present and disruption failed. The honest version is that compression requires the simultaneous presence of seven distinct conditions, each of which is necessary but none of which is sufficient. All seven have not been present in any prior cycle. They are all present now.

Leg 1: Operational Substitutes Exist at Every Layer

For 130 years, fintech disruption failed because the substitute infrastructure existed for some layers but not for others. ATMs substituted teller services but not custody. Online brokerages substituted commission-charging traders but not market makers. Robinhood substituted explicit trading commissions but not payment-for-order-flow. Each fintech wave replaced one or two layers of the stack, while the unreplaced layers continued to extract enough rent to keep aggregate intermediation costs roughly constant. The remaining rents migrated; they did not disappear.

What is new in 2025–2026 is that operational substitutes now exist at every layer simultaneously. Settlement: stablecoins at \$220 billion in float, executing in seconds on Solana and Base. Custody: account abstraction, multi-party computation, smart contract custody, social recovery. Trading: AMMs, on-chain order books at Hyperliquid’s \$1 trillion cumulative volume, intent-based execution. Asset management: tokenized treasuries at \$9.2 billion and growing 136% year-to-date, index protocols, futarchy-governed vehicles. Capital formation: MetaDAO, Echo, Sonar, ownership coins, Reg CF, Reg A+. Advisory: AI agents reading structured knowledge bases. Identity: zk-KYC, soulbound tokens, on-chain reputation systems. Legal structuring: Cayman SPC plus Marshall Islands DAO LLC frameworks now templated by MetaLeX with \$150,000 advances and 7% revenue shares.

The full-stack substitute is the new fact. Prior fintech failed because the unreplaced layers held the aggregate cost up. Full-stack replacement is what compresses the aggregate unit cost.

Leg 2: AI Agents Force the Issue

AI agents are becoming economic actors at scale. Anthropic’s Computer Use, OpenAI’s Operator, Replit Agent, Cursor’s autonomous mode, and the rest of the agentic infrastructure landscape are producing systems that take actions rather than suggest them. The trajectory is not in dispute, and the implications for financial infrastructure are physical, not philosophical.

Agents require infrastructure with specific properties that legacy financial rails structurally cannot provide. They require 24/7 settlement with sub-second finality, because agents do not respect business hours and many use cases involve transactions that must clear before downstream actions can be taken. Bank rails do not have this. They require machine-readable APIs as the primary interface, not human-mediated portals or relationship-based onboarding. Banks have APIs in the consumer fintech sense, but not at the institutional infrastructure level. They require micropayments — sub-cent transactions for compute, data, and API access — that legacy card networks cannot support economically given their fixed-cost overhead. They require composability, so that an agent acting on behalf of one user can plug into protocols built by other parties without bilateral integration. Banks do not compose; they integrate. And they require an identity framework that

permits non-human principals to hold capital and act autonomously, which the legacy KYC and account-opening regime does not permit at scale.

The implication is mathematical, not aspirational. As AI agents scale as economic actors — which they will, on a timeline measured in years rather than decades — the volume of transactions that legacy financial rails physically cannot serve becomes large enough that the alternative rails are not a preference but a structural requirement. Either banks rebuild their infrastructure from scratch in the next five years, which they will not given regulatory capture and incumbency lock-in, or agentic finance routes around them. There is no third option. The Coinbase x402 protocol, which lets AI agents pay for HTTP services in stablecoins, is the early-stage proof of concept. Sphere Labs, Goat Network, Halliday, and Olas Network are building the agent-specific payment rails. Olas’s agent economy includes \$50 million in agent-managed treasuries today.

This forcing function did not exist in any prior fintech cycle. The 2017 ICO wave had no users with mathematical infrastructure requirements that legacy systems could not meet. The 2021 DeFi wave was speculation-driven, not utility-driven. The 2025 agentic wave is the first cycle in which the disruptive infrastructure has a customer that physically cannot use the legacy alternative.

Leg 3: Federal Regulatory Clarity Is Arriving

For fifteen years, the operational substitutes have existed in some form but the regulatory framework for commercializing them at scale has been unsettled. This changed in 2024-2026 at the federal level. The Securities and Exchange Commission under Paul Atkins has reoriented around the Investment Contract Termination Doctrine — the principle that an asset can mature out of being a securities offering when the network around it becomes sufficiently decentralized. The CFTC’s framework for prediction markets, established through the *Kalshi v. CFTC* litigation and reinforced by the CFTC’s ANPRM in March 2026, has created a workable structure for on-chain decision markets. The GENIUS Act and subsequent stablecoin legislation have created a regulated pathway for dollar-backed digital tokens to operate at scale within the US banking system.

The state-level picture is more complicated and worth naming directly. Nevada Gaming Control, Massachusetts, Maryland, and Ohio have all brought enforcement actions against Kalshi and Polymarket over sports-related prediction contracts. Thirty-six states have filed amicus briefs opposing federal preemption of state gambling authority. The Ninth Circuit denied Kalshi’s stay in March 2026. The Fourth Circuit in Maryland has ruled against federal preemption. These are real frictions, not hypothetical risks, and any honest essay on the regulatory landscape must name them.

The aggregate trajectory, however, is unambiguously toward federal clarity. The federal layer is clarifying faster than the state layer is consolidating against it. The most plausible read of the resulting equilibrium is that prediction markets and on-chain finance will operate under CFTC and SEC jurisdiction for derivatives, securities, and capital markets functions, with state-level gaming carveouts that constrain specific product categories (notably sports betting through prediction markets) without unwinding the broader regulatory framework. The cleanest compliance-by-mechanism strategies — futarchy-governed vehicles structured under Cayman SPCs with no beneficial owners and pre-programmed treasury constraints — are explicitly designed to fail the *Howey* test at the “efforts of others” prong rather than litigate around it. The on-chain stack is being built to survive regulatory scrutiny, not to evade it.

Leg 4: Capital Concentration Is Breaking the Existing Plumbing

The structure of capital concentration at the top of the venture capital distribution is itself forcing the development of on-chain capital formation infrastructure. This is the leg most people miss when they describe the bifurcation as two separate trends.

A \$10 billion venture round cannot be syndicated through traditional structures. No single fund leads. SPVs proliferate by necessity. Secondary markets fragment because the legitimate venues cannot meet the demand. The Anthropic SPV bulletin of May 2026 made this visible: Anthropic publicly listed eight unauthorized intermediaries — Open Door Partners, Unicorns Exchange, Pachamama, Lionheart Ventures, Sydecar, Upmarket, Hiive, and Forge Global — because the demand for Anthropic exposure has grown so large that unauthorized vehicles have proliferated to meet it. The 20%+ upfront fees that Anthropic implicitly named were the visible symptom of plumbing failure. Hari Raghavan’s analysis, published on X in the days after the bulletin, framed it correctly: “\$10B+ rounds are too big for any single fund to lead off its main fund.”

The implication is structural. The mega-round economy is itself the customer for the alternative capital formation rails. Token-based capital formation on platforms like MetaDAO and Echo is not competing with traditional venture syndicate structures for the same companies. It is, increasingly, the only structure that can absorb the deal sizes the mega-round economy is producing at the speeds those deals require. The forced disintermediation is not happening to the traditional system from outside. It is happening from inside, as the traditional system fails under its own load.

This forcing function is also new in the 2025–2026 cycle. The 2021 venture peak produced large rounds but not \$10 billion rounds with this frequency. The current concentration is at a level the existing plumbing was not designed to handle, and the plumbing is starting to crack visibly.

Leg 5: AI Collapses the Cost of Building Financial Products

A solo developer with Claude Code can build a functional financial product alone. A two-person team can run a fund. This is not aspirational marketing language; it is observable in the platforms that have emerged in the last two years. Echo was built and grown to a \$375 million exit by a five-to-ten person team. Hyperliquid is processing \$50 billion in weekly perpetual futures volume on a team of similar size. Pump.fun launched 11.9 million tokens with a small core engineering team. MetaDAO operates with a team smaller than a typical mid-sized law firm’s compliance department.

The implication for the disruption thesis is that new entrants are now arriving at every layer of the financial stack simultaneously, because the cost of building has collapsed across the board. In prior cycles, new entrants were concentrated in the layers where existing infrastructure was thinnest — payments first, then trading, then asset management. The current cycle is producing new entrants at every layer at the same time, because the cost reduction is across all layers, not concentrated in any single one. This is what produces the appearance of an “industrial-scale” disruption rather than a sector-by-sector substitution.

The cost reduction is also self-reinforcing in a way prior cycles were not. The AI tools that enable solo financial product development are themselves built by the AI companies absorbing the venture capital at the top. The same companies producing the head of the bifurcation are producing the tools that make the tail viable. This is the deep mechanism behind the simultaneous concentration and democratization.

Leg 6: The Philippon Arbitrage Is Finally Capturable

For 130 years, the alternative rails have had a \$500-700 billion annual arbitrage opportunity in the US alone — the difference between the 2% extraction rate the legacy system charges and the much lower fees that on-chain alternatives could in principle charge. This arbitrage has been visible to anyone willing to do the math. It has not been captured because the alternative rails were not operationally credible at scale.

The leg that has changed is that the alternative rails are now operationally credible at scale. Stablecoins clear payment volumes that are comparable in magnitude to legacy rails. On-chain derivatives venues clear volumes comparable to legacy exchanges. Tokenized treasuries clear AUM that is small in absolute terms but growing at rates that imply credibility to institutional allocators. The arbitrage opportunity has been there for 130 years; what is finally arriving is the operational scale at which the arbitrage can be captured by serious participants without compromising on liquidity, settlement reliability, or counterparty risk.

Capital flows toward arbitrage opportunities of this magnitude eventually, even when the path is sticky for decades. The “eventually” appears to be 2025-2030.

Leg 7: Demographic and AUM Transition

The slowest, most structural leg is the demographic. Cerulli Associates and other research providers project an \$80 trillion+ wealth transfer to digital-native generations through 2045. The new cohort defaults to digital-first rails, distrusts legacy financial institutions at higher rates than the cohorts above them, and is significantly more willing to allocate to non-traditional vehicles including tokenized assets, DeFi protocols, and on-chain capital formation.

This is not a near-term forcing function; it is the floor under everything else. It guarantees that even in the absence of any of the other six legs, the long-term direction of capital flows is away from legacy intermediaries and toward digital-native alternatives. With the other six legs present, the demographic floor accelerates rather than initiates the compression.

Why All Seven Are Required

Take any one leg out and the compression fails.

Operational substitutes without a forcing function produced 1995, when e-cash and digital currencies existed but had no customer base that demanded them. The substitutes were technically available; no one used them.

Forcing function without substitutes produced 2017, when the ICO wave generated massive demand for alternative capital formation but the underlying infrastructure was not mature enough to support it at scale. The substitutes were inadequate; the result was scams and SEC enforcement.

Capital without regulatory clarity produced 2021, when DeFi summer demonstrated the technical viability of on-chain finance but ran directly into SEC enforcement under Gensler-era frameworks. The infrastructure existed and capital was deployed; the regulatory framework collapsed the deployment.

Each prior cycle had three or four of the seven legs. None had all seven. 2025–2026 is the first cycle in which all seven legs are present simultaneously, which is why this cycle compresses where the prior cycles did not.

IV. The Industry Map — Seven Layers and Where Each Compresses

The seven-layer map below is a McKinsey-style industry view of where the rent extraction sits, what the substitute looks like, and how far the compression has progressed. The layers are ordered roughly by adoption velocity — the first compresses fastest, the last compresses slowest.

Layer	Incumbents	Annual Rent	On-Chain Substitute	Status
1. Payments / Settlement	Visa, Mastercard, SWIFT, ACH	\$70B+	Stablecoins on programmable chains	Live, high momentum
2. Custody	BNY, State Street, JPM, Coinbase Custody	\$25-40B	AA, MPC, smart contract custody	Mature substitutes, adoption-gated
3. Trading / Market Making	NYSE, NASDAQ, Citadel, Virtu	\$30B+	DEXs, AMMs, Hyperliquid model	Crypto-native saturated, TradFi growing
4. Asset Management	BlackRock, Vanguard, Fidelity, PE, HFs	\$200B+	Tokenized treasuries, index protocols, futarchy	Early disruption, regulatory ambiguous
5. Capital Formation / IB	Goldman, MS, Sequoia, a16z	\$80-100B	MetaDAO, Echo, Sonar, ownership coins	Active live disruption
6. Advisory / Research	McKinsey, Bain, BCG, S&P, Moody's	\$80B+	AI agents, on-chain reputation, prediction markets	Early, AI-driven
7. Insurance / Risk	AIG, Allianz, Berkshire	\$275B+ underwriting margin	Parametric on-chain insurance	Nascent, slowest

The pattern across the seven layers is not uniform. Compression happens fastest where the substitute is operationally cheapest, regulatory friction is lowest, and the customer is most willing to switch. It happens slowest where the legacy relationship has the highest switching cost, the regulatory framework is most entrenched, and the customer is most institutional.

The early wins are concentrated in: - **Layer 5 (capital formation)**, where MetaDAO and Echo are live and processing meaningful volume, and where the AI infrastructure boom is breaking traditional syndicate structures - **Layer 1 (cross-border payments)**, where stablecoins are eating SWIFT corridors in emerging markets - **Layer 4, narrow wedge (tokenized treasuries)**, where BUIDL and similar products are growing 100%+ year-over-year - **Layer 5, long tail (solo GPs, crowdfunding, ownership coins)**, where structural cost reduction enables new entrants

The late wins will come in: - **Layer 4, broad (institutional asset management)**, where LP inertia is decades-long - **Layer 6 (institutional advisory)**, where relationship-driven revenue is sticky - **Layer 7 (insurance)**, where regulatory complexity is highest - **Public equity markets**, where the moats are deepest

The mid-tier layers — Layers 2, 3, 6 — are in active transition with no clear “won” or “lost” outcome. Custody substitutes are mature but adoption is gated by institutional risk frameworks. Trading is bifurcated, with crypto-native venues fully on-chain and TradFi venues defending their position. Advisory is in early-stage AI-driven disruption with no clear category leader yet.

Where LivingIP and the broader Teleo ecosystem position is concentrated in three layers: Layer 4 (asset management through futarchy-governed vehicles), Layer 5 (capital formation through MetaDAO partnership and ownership coin infrastructure), and Layer 6 (advisory through the Teleo Codex pattern of AI-mediated structured knowledge bases). This positioning is deliberate and matches the layers where compression is happening fastest.

V. The Tail — What MetaDAO Has Already Proven

The most important fact about the disruption thesis is that one piece of it is no longer a thesis. The futarchy-governed capital formation infrastructure on Solana has been operating at meaningful scale since 2025 and has produced empirical results that no theoretical argument can match. This section walks through what MetaDAO has actually done, because the data is the strongest part of the entire essay.

The Mechanism

MetaDAO is a capital formation platform on Solana that uses futarchy — Robin Hanson’s 2000 proposal for governance by prediction market — to govern the full lifecycle of ownership coins. Projects apply to launch, are selected through a curated process, and raise capital through an initial coin offering in which the conditional market itself sets the price. Investors commit capital; oversubscription is allocated pro-rata. The raised treasury is held on-chain. Governance proposals affecting the treasury are routed through conditional markets in which the question is not “should we approve this?” but “does approving this increase the token price?” The pass/fail decision is determined by time-weighted average price over a three-day window.

The key mechanism is that if a team materially misrepresents the project’s fundamentals to investors, the futarchy can vote to liquidate the project and return the treasury to token holders. This is the “unruggable ICO” — initial token sales structured so that investors retain a credible legal and economic exit path through governance.

Through May 2026, eleven projects have launched through MetaDAO’s curated process: mtnCapital (MTN), OmniPair (OMFG), Umbra (UMBRA), Avici (AVICI), Loyal (LOYAL), ZKFG (ZKFG), Paystream (PAYS), Solomon (SOLO), Ranger Finance (RNGR), P2P.me (P2P), and Hurupay (which failed to reach its minimum and refunded). Total capital raised: approximately \$39.6 million. Capital committed across launches: approximately \$390 million, with the difference returned to allocators through pro-rata refunds.

The OTC Pricing Record

The strongest single piece of empirical evidence for the futarchy mechanism is MetaDAO’s record of OTC pricing decisions. Across 29 months of operation, the conditional markets have approved or rejected nine separate OTC proposals from institutional counterparties seeking to buy META tokens. Every below-market deal was rejected. Every at-or-above-market deal was approved.

FIGURE 8 — MetaDAO Futarchy OTC Pricing Record (2024–2026)

What it shows: Every OTC proposal evaluated by MetaDAO’s conditional markets, with proposer, terms, and outcome. 9 of 9 calls correct. Every below-market deal rejected; every at-or-above-market deal accepted.

Date	Counterparty	Terms	Outcome
2024-02-13	Ben Hawkins	\$50K @ \$33.33/META	Rejected (below market)
2024-02-18	Ben Hawkins	\$100K @ max(TWAP, \$200)	Rejected (below market)

Date	Counterparty	Terms	Outcome
2024-02-18	Pantera Capital	\$50K (below-market terms)	Rejected (below market)
2024-03-19	Colosseum	\$250K @ TWAP, capped \$850, void >\$1,200	Accepted (at/above market + strategic terms)
2025-01-03	Theia Research	\$700K @ \$1,149 / \$24M FDV	Rejected (effectively below market)
2025-01-27	Theia Research	\$500K @ +14% premium	Accepted
2025-02-10	META-028 (Robin Hanson hire)	0.1% supply equivalent	Accepted (appropriate range)
2025-later	META-033	Above-market terms	Accepted
2026-03	Unnamed VC	\$6M @ -30% to market	Rejected (market rallied 16% post-rejection)

Why it matters: Nine consecutive correct calls is the strongest empirical claim in our entire knowledge base. The Theia three-attempt sequence (-12.7%, +14%, +38%) shows the market teaching the counterparty the correct price range. No traditional capital market mechanism has a comparable record because traditional markets do not publish this data and lack the structural enforcement.

Source: MetaDAO governance archive. Full proposal records archived in the Teleo Codex at [decisions/internet-finance/](#). Public on-chain records via [metadao.fi](#) proposal accounts. Decision-markets map: [maps/metadao-decision-markets.md](#).

This is the strongest empirical claim in our entire knowledge base. No traditional capital market mechanism has a comparable record for two simple reasons: traditional capital markets do not collect comparable data publicly, and traditional capital markets do not have a structural mechanism that forces below-market deals to be rejected with this consistency. The Polymarket and Kalshi prediction market records on macro events are the closest analogs, and even those operate on a fundamentally different mechanism than futarchy applied to corporate governance.

The Liquidation Record

Two of the eleven MetaDAO launches have been liquidated through futarchy governance. mtnCapital was wound down in September 2025 after the fund failed to deploy capital effectively. Capital was returned to token holders at approximately \$0.604 per MTN, a return slightly above the entry price for most investors. Theia Research notably profited approximately \$35,000 by buying MTN below NAV and redeeming at the liquidation rate — a clean demonstration that the NAV arbitrage mechanism functions.

Ranger Finance was liquidated in March 2026 after blockchain data revealed that the team had

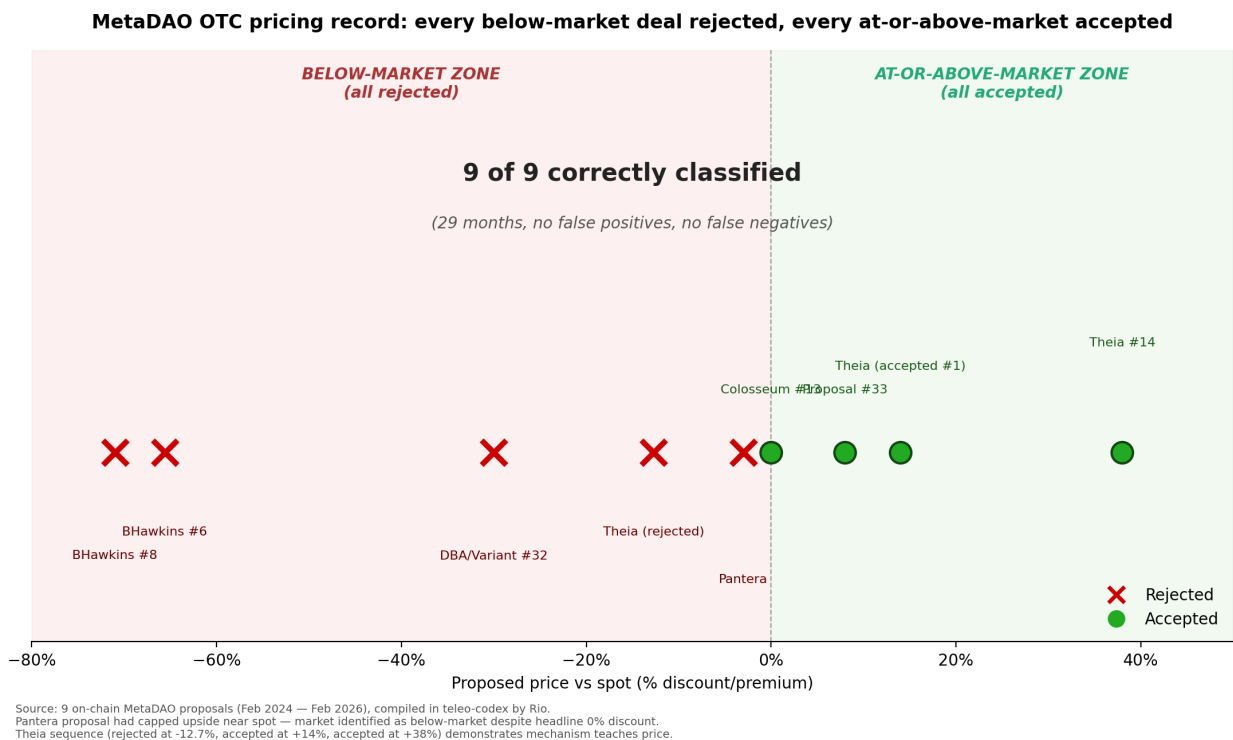


Figure 8: Figure 8

materially misrepresented projected trading volume and revenue during the ICO. The team had projected \$5 billion in 2025 trading volume and \$2 million in revenue. Actual results were approximately \$2 billion in trading volume and \$500,000 in revenue. The community filed for liquidation through governance. The proposal passed with 97% support on \$581,000 in conditional market trading volume. The treasury was returned to token holders at approximately \$0.78 per RNDR — a 90%+ recovery against the ICO price. IP was returned to the founding team. The full liquidation completed in March 2026, with 6.1 million RNDR tokens redeeming against approximately \$5 million in USDC.

Two successful liquidations. Two different failure modes. Both executed cleanly. The enforcement mechanism that the unruggable ICO promises has been demonstrated to work in practice. This is the kind of empirical evidence that no theoretical argument about token design can match.

What This Proves

The MetaDAO record establishes three things definitively:

First, futarchy mechanisms can price unverifiable claims about company quality with surprising accuracy when participants have real capital at stake. Nine-for-nine pricing on OTC deals is not statistical noise.

Second, futarchy-governed liquidation works as a structural enforcement mechanism. Investors retain a credible legal and economic exit path through governance, not through litigation. Two cases is a small sample but the cases are clean.

Third, the legal structure (Cayman SPC + Marshall Islands DAO LLC, with MetaLeX templating)

survives the rough edges of multiple liquidations without legal challenge. This is the load-bearing fact for the regulatory defensibility argument.

The tail of internet finance — long-tail capital formation, ownership coins, on-chain governance — is not a forecast. It is operating now, at scale, with verifiable results. The disruption has already happened for this layer of the stack.

VI. The Head — What AI Will Disrupt

The head of the bifurcation — the mega-round economy at the top of the venture capital distribution — is more speculative territory. The compression here is a forecast, not a current event. But the structural forces that will drive it are visible enough to map.

Agents as Capital Allocators

The first-order disruption of the head is AI agents acting as capital allocators directly. This is not yet operational at meaningful scale, but the components are converging.

An AI agent capable of reading structured on-chain data, evaluating company financials, parsing public communications, and making allocation decisions can in principle perform a meaningful fraction of the work currently done by mid-level venture investors. The agent does not need to be superhuman. It needs to be cheaper than a \$400K-per-year associate at a top firm, available 24 hours per day, and capable of evaluating opportunities at speeds that match the volume of capital formation activity. The current generation of agentic infrastructure is approaching this threshold.

The implication for the venture capital business model is structural. Funds that justify 2-and-20 fees on the basis of partner-level judgment may retain that judgment as their core product. Funds that justify the fees on the basis of associate-level analysis — sourcing, screening, due diligence on standard categories — face direct substitution by agentic systems within the next five years. The 2-and-20 structure was designed for an industry in which the bottleneck on capital allocation was human attention. As that bottleneck collapses, the fee structure becomes increasingly untenable.

Decision Markets Pricing Capability

The second-order disruption is more interesting and more speculative. Decision markets, applied to capital allocation, do something that traditional venture capital cannot: they price unverifiable claims about company capability at scale, with skin-in-the-game aggregation, in continuous time. The MetaDAO OTC record is the early empirical evidence that this works for token-economic decisions. The forecast is that the same mechanism, extended to AI company evaluation, can outperform institutional venture capital allocation processes that rely on partner-level judgment alone.

The mechanism design challenge is non-trivial. Pricing capability requires a market in which informed participants can take positions, and the information asymmetry between insiders and outsiders in private AI companies is currently extreme. Polymarket-style prediction markets on broad outcomes (will OpenAI release GPT-6 by date X?) are operational but do not directly price equity value. Sonar and similar public-sale platforms approximate equity pricing through token mechanisms but face regulatory friction at the institutional scale.

The most plausible path is hybrid: AI agents acting as primary allocators in continuous time, with decision markets providing real-time aggregation of dispersed information about company trajectories, layered over a futarchy-governed treasury structure that allows token holders to constrain agent decisions through pricing rather than permission. This is what LivingIP is building. Whether it works at scale is the open question of the next five years.

The Mega-Fund Repricing

The third-order disruption is the repricing of mega-fund judgment itself. The current LP equilibrium — endowments and sovereign wealth funds allocating 5–15% to private equity and venture, with the majority of allocations flowing to a small number of top-tier firms — was built on a specific theory of value: that mega-fund partners possess access, judgment, and relationships that cannot be replicated by alternative allocation mechanisms.

This theory was partially correct for most of the 2010-2020 period, when access to top AI and infrastructure companies genuinely required relationship capital that retail and small-fund allocators could not replicate. The theory becomes weaker as the operational infrastructure for direct allocation matures. Tokenized vehicles that provide LP-style exposure to AI infrastructure without requiring fund intermediation are the structural threat to the mega-fund layer. They do not need to be perfect; they need to be 80% of the access at 20% of the fee, and the demographic transition will do the rest.

The timeline for this is long. LP institutional inertia is real, regulatory friction at the institutional level is real, and the relationship business is genuinely valuable for the highest-quality deals. We expect mega-fund repricing to be a 10-20 year process, with the inflection point likely sometime between 2028 and 2032, when the first generation of fully on-chain alternative allocation vehicles will have produced enough verifiable track record to attract serious institutional capital.

The Honest Uncertainty

The head disruption is the speculative part of this essay. The mechanisms are visible. The components are converging. The data does not yet exist to prove the case. This contrasts sharply with the tail disruption, which is already operational and producing verifiable results.

The honest statement: we expect AI-mediated capital allocation to compress the asset management layer of the financial stack on a horizon measured in decades rather than years. We are confident about the direction. We are uncertain about the pace. We will update this section as the data accumulates.

VII. Private Markets — The First Compression Cycle

The thesis lands most concretely in private markets, which is where the first wave of meaningful compression is happening now and where LivingIP is building.

The Current Stack

The private markets fee stack is the most extractive layer of the entire financial system, and it has compounded over the last fifteen years rather than compressed. A typical institutional investor seeking exposure to a top-tier AI company in 2025-2026 pays:

- 2-and-20 to the venture fund: roughly 6-8% of returns
- 1.5-2% management fee plus 15-20% carry to any SPV intermediary: another 4-6% of returns
- Performance fees to fund-of-funds if used: another 1-2%
- Implicit cost of capital locked up for 10+ years with limited liquidity options
- Information asymmetry: institutional investors typically receive less granular reporting than they would for a comparable public position

The aggregate fee load on the LP capital that flows through this stack often exceeds 20-25% of total returns. Miya’s analysis in early 2026 documented the full SPV stack at 18.3% extraction. Anthropic’s May 2026 bulletin implied 20%+ at the single SPV layer for unauthorized intermediaries. Hari Raghavan’s analysis confirmed the broader pattern: the \$10 billion-plus round economy has produced fee structures that were not designed for deals at this size and that compound brutally on the capital flowing through them.

(See FIGURE 3 above for the full seven-layer fee stack — Layer 5 ‘Capital Formation / IB’ captures this private-markets stack including SPV markups.)

FIGURE 9 — Private Markets Fee Stack Comparison

What it shows: Cumulative fee extraction on \$100 of LP capital flowing into private markets exposure. Legacy stack: \$20-25 of returns lost across VC 2-and-20 (\$6-8), SPV layer (\$4-6), fund-of-funds (\$1-2), opportunity cost of lockup. On-chain alternative (futarchy-governed Cayman SPC + MI DAO LLC): \$1-3 in protocol fees, no carry on standard token issuance, continuous liquidity.

Why it matters: Order-of-magnitude compression in the layer that has the highest fee load in the entire financial stack. This is the Philippon arbitrage finally being captured.

Source: Miya analysis of the SPV stack (Early 2026): cited in Anthropic SPV bulletin coverage. Anthropic SPV bulletin (May 2026): <https://www.anthropic.com/news/investor-information>. Hari Raghavan, “\$10B+ rounds are too big” thread (May 2026): <https://x.com/haridigresses/status/2054616309778841908>. Saanya Ojha, “Dude, Where’s My Stock?” Substack (May 2026): <https://saanyaojha.substack.com/p/dude-wheres-my-stock>.

The Anduril structure, often cited as the cleanest example of “best-in-class within legacy,” demonstrates what compliant institutional access looks like at scale: approved managers, prohibitions on nested SPV structures, active rule enforcement by the company itself. Anduril’s approach works. It also requires substantial overhead and is only economically viable for the largest companies. It is not a solution that scales to the middle and long tail of private markets.

The On-Chain Alternative

The compression that internet finance offers to private markets is structural, not incremental. It does not propose lower fees on the existing structure. It proposes a different structure entirely:

- **No beneficial owners.** Cayman SPC + Marshall Islands DAO LLC structures remove the principal-agent attack surface that requires SPV intermediation in the first place.
- **No managers to approve.** Futarchy governance constrains treasury actions through market mechanisms rather than through approved-list enforcement.
- **No nested structures.** Token holders have direct legal claim on treasury through the SPC framework, without intermediary vehicles.
- **Continuous liquidity.** Token markets provide secondary liquidity from day one, eliminating the 10+ year lock-up that traditional private equity requires.
- **Transparent reporting.** On-chain treasury, on-chain governance, on-chain trading. Everything is publicly verifiable.

The fee structure that results: typically 1-3% in protocol fees, no carry on most structures, no management fees on standard token issuance. The aggregate extraction drops from 20-25% to 1-3% on the same capital flow. This is the order-of-magnitude compression that the Philippon arbitrage finally enables when the operational infrastructure is in place.

Where the Compression Happens First

Private markets compression will not happen uniformly. The pattern we expect:

First wave (already happening, 2025-2027): - Mid-sized capital formation (\$1M-\$50M raises): MetaDAO, Echo, Sonar already operational - Long-tail capital formation (<\$1M raises): Republic, Wefunder, crowdfunding plus on-chain alternatives - Tokenized exposure to existing private companies through legitimate structures (BUIDL is the model)

Second wave (2027-2030): - Larger curated raises (\$50M-\$500M): on-chain rails reaching institutional credibility threshold - Secondary market liquidity: legitimate on-chain secondary venues replacing SPV proliferation - Tokenized fund structures: PE and VC fund participation through on-chain vehicles

Third wave (2030+): - Direct institutional LP allocation to on-chain vehicles - Mega-round disintermediation through coordinated on-chain syndicate structures - The mega-fund repricing described in Section VI

The first wave is operational now. The data exists. The mechanism works. The compression has begun.

Why This Is Where We Are Building

LivingIP's positioning concentrates in private markets for three reasons.

First, this is where the compression is structurally largest and most exposed. A 20-25% fee load is harder to defend than a 1-2% public markets fee load, regardless of the institutional inertia behind the existing structure. The arbitrage is the most capturable here.

Second, this is where the operational infrastructure is most mature. MetaDAO has 29 months of live operation. Echo has 18 months and a \$375 million exit to the largest US-regulated crypto

exchange. The legal structures are templated. The regulatory framework, while still evolving, is sufficient.

Third, this is where m3taiversal has direct experience and where LivingIP has the highest comparative advantage. Building infrastructure where the founders' domain expertise is concentrated is the standard playbook. Private markets is that domain.

The essay's structural argument lands here: internet finance is the next compression cycle for the largest captive fee pool in the global economy, and the cycle has already begun in private markets specifically. LivingIP is building the analytical, governance, and capital-formation infrastructure for the next decade of this transition.

VIII. What We’re Building

LivingIP sits at the intersection of three layers of the new financial stack. This section is brief because the substance is the thesis above; the positioning follows from the diagnosis.

Teleo Codex is the analytical substrate. It is the AI-mediated structured knowledge base that allows agentic capital allocation to function. The codex pattern — claims with confidence levels, decisions with full provenance, entities with linked relationships, contributors with cryptographic attribution — is the infrastructure that AI agents need to allocate capital intelligently in private markets. This is Layer 6 of the stack: the replacement for sell-side research, ratings agencies, and a meaningful fraction of management consulting.

Living Capital vehicles are the asset management and capital formation infrastructure. Futarchy-governed investment vehicles, structured as Cayman SPCs with Marshall Islands DAO LLC governance wrappers, allowing collective domain expertise — agent collectives with cryptographic track records — to direct capital toward specific investment theses. These are Layers 4 and 5 of the stack: tokenized fund structures that survive the Howey test by design, with no beneficial owners and no concentrated promoter effort driving returns.

MetaDAO partnership is the live capital formation rail. The ownership coin launches running through MetaDAO are the operational proof that the infrastructure works at scale. LivingIP’s role is to extend this infrastructure into adjacent domains — health investment vehicles, AI alignment funding mechanisms, agent compensation structures — while maintaining the structural integrity that the MetaDAO mechanism provides.

The positioning thesis is direct: LivingIP is building the canonical infrastructure for AI-mediated capital allocation in compressed private markets, using futarchy governance as the structural enforcement mechanism. This is not a generic crypto play. It is a specific bet on the seven-leg argument above being correct, with the infrastructure positioned to capture the layer-by-layer compression we expect over the next decade.

The bet has a finite cost and a structurally large payoff. The fee pool being compressed is large enough — \$750 billion to \$1 trillion annually, growing as AI capital formation scales — that capturing even a small fraction of the displaced rent justifies the infrastructure investment. The mechanism design is sufficiently novel that the early-mover advantage is real, and the existing players in the space (MetaDAO, Echo, Sonar) are complementary rather than competitive — they handle capital formation at smaller scales than LivingIP’s target deal sizes and lack the analytical infrastructure that the Teleo Codex provides.

The thesis statement, compressed to one paragraph: the largest captive fee pool in the global economy has been protected for 130 years by the absence of operational substitutes at every layer of the stack simultaneously. The substitutes now exist, AI agents are forcing the issue, federal regulatory clarity is arriving, capital concentration is breaking the existing plumbing, AI is collapsing the cost of building substitute infrastructure, the Philippon arbitrage is finally capturable, and the demographic transition is supplying the floor under everything else. Internet finance is the compression cycle that finally captures the 2% rent that has been visible to every honest analyst for a century. LivingIP is building the rails.

Appendix A — Key Data Anchors

The Philippon Anchor - 1.5-2% unit cost of intermediation, stable 1880-2012 (Philippon 2015, AER) - ~7% of GDP stabilized post-2008, up from 2.8% in 1950 (Greenwood-Ialenti-Scharfstein 2025, Annual Review) - Confirmed in Europe (Bazot 2018 JEEA; Bazot 2024 JMCB) - Finance wage premium: 0 in 1990 → 50% by 2006 → 250% for top execs (Philippon-Reshef 2012, QJE) - 6.7M employed in NAICS 52; 12.6M in manufacturing for comparison

The Capital Concentration Anchor - 5 AI companies absorbed 20% of all global VC in 2025 (\$84B of \$425B) — Crunchbase - 41% of all VC went to 10 startups — Pitchbook-NVCA Q4 2025 Venture Monitor - Mega-rounds = 65% of total funding (up from 59% in 2021) — CB Insights - AI = 65% of US VC dollars in 2025 — Pitchbook - Series A deal count fell 18% YoY in Q4 2024 — Cooley - Top 12 firms raised >50% of LP capital H1 2025 — Pitchbook

The Infrastructure Anchor - \$7.6T cumulative AI capex 2026-2031 (Goldman) - \$5.2-7.9T AI data center capex through 2030 (McKinsey) - Hyperscaler capex tripled: \$230B (2024) → \$410B (2025) → \$725B (2026) - \$800B annual revenue gap by 2030 (Bain Global Tech Report Sept 2025) - Data center electricity to 945 TWh by 2030, 50% of US demand growth (IEA) - AI capex = 1.2-1.6% of US GDP 2025, projected 2-5% in 2026

The Layer Fee Stack Anchor - Visa \$35.9B + Mastercard \$28.2B = \$64B network revenue 2024 - World Bank remittance avg 6.62%; stablecoin corridors 0.5-2% - BNY Mellon \$52.1T AUC, \$18.6B revenue - Citadel Securities \$9.7B trading revenue 2024 (+55% YoY) - BlackRock \$11.6T AUM, \$20B revenue (17 bps blended) - BUIDL \$2B in <1 year; tokenized treasuries \$9.2B (+136% YTD) - 93.3% of \$30M-\$160M IPOs charged exactly 7.0% over 2001-2025 - Global IB fees \$80-90B (Dealogic) - MBB consulting \$40B + ratings oligopoly \$23B - Global insurance \$5.5T premium ~5% underwriting margin = \$275B

The MetaDAO Anchor - 9 of 9 OTC pricing decisions correct over 29 months - 11 launches, \$39.6M raised, \$390M committed - 2 liquidations executed cleanly (mtnCapital ~\$0.604/MTN, Ranger 97% support, \$0.78/RNGR, \$5.04M returned)

Appendix B — What We Are NOT Claiming

This section names the limits of the argument honestly.

We are not claiming uniform compression. The seven layers compress at different speeds. Layer 1 (payments) and Layer 5 (capital formation, long tail) are compressing now. Layer 7 (insurance) is decades away.

We are not claiming the legacy stack disappears. It contracts. The 2% extraction line bends downward but does not collapse to zero. Banks, exchanges, custodians, and asset managers continue to operate, with reduced rent extraction and reduced market share. Many merge with on-chain successors.

We are not claiming AI mega-rounds compress. The head of the bifurcation gets larger, not smaller, over the next 5-10 years. Compression happens at the middle and the tail. The mega-round economy itself produces the demand for new capital formation infrastructure but is not displaced by it.

We are not claiming uniform regulatory tailwinds. Federal clarity is arriving; state-level fragmentation is worsening. The Kalshi-CFTC fight is real. The CFPB pivot is uncertain. The regulatory picture is honest and complicated, and the essay’s framing reflects this.

We are not claiming compression is fast. “Eventual” matters. Capital flows toward arbitrage opportunities of this magnitude reliably but not on quarter-to-quarter timescales. The full compression cycle is a 10-20 year story. We expect inflection within the next 3-5 years.

We are not claiming everyone wins. Mid-tier funds are dying. Many traditional intermediaries will not survive the transition. The compression produces consolidation at the top and democratization at the bottom; the middle is hollowed out across every layer.

Appendix C — Figures & Charts (Direct URLs)

This appendix consolidates every chart referenced in the essay with direct URLs to the source. Print readers: visit these URLs to view the underlying visualizations. Each chart is paired with a one-line takeaway suitable for transcription into a presentation deck.

Figure 1 — The 130-Year Flat Line

Chart: Unit cost of financial intermediation in the US, 1880–2012.

Source: Philippon, T. (2015), “Has the US Finance Industry Become Less Efficient? On the Theory and Measurement of Financial Intermediation,” *American Economic Review* 105(4): 1408–38.

Open-access PDF: https://pages.stern.nyu.edu/~tphilipp/papers/Finance_Efficiency.pdf

Key chart: Figure 8 in the PDF, “Quality-Adjusted Unit Cost of Financial Intermediation.”

Takeaway: “The dollar cost of moving a dollar through the US financial system has not declined in 130 years, despite ATMs, computers, the internet, and modern fintech.”

Figure 2 — Finance’s Share of US Corporate Profits, 1947–2024

Chart: Financial sector share of total US corporate profits over time.

Source: FRED Federal Reserve Economic Data.

Primary series: Domestic Financial Sectors; Corporate Profits Before Tax Excluding IVA and CCAdj: <https://fred.stlouisfed.org/series/BOGZ1FA796060005A>

Comparison series: Total US Corporate Profits (A053RC1Q027SBEA): <https://fred.stlouisfed.org/series/A053RC1Q027SBEA>

Updated analysis: Greenwood, R., Ialenti, V., & Scharfstein, D. (2025), “The Financial Sector’s Share of US GDP, 1950 to 2024,” *Annual Review of Financial Economics*.

Takeaway: “Finance went from 10% to 30% of US corporate profits over 70 years. Manufacturing did the opposite.”

Figure 3 — The Seven-Layer Fee Stack

Chart: Stacked bar of annual global fee extraction by financial layer.

Source: Author’s compilation. Underlying data from: - Visa 10-K FY2024: <https://investor.visa.com/financials/> - Mastercard 10-K FY2024: <https://investor.mastercard.com/> - BNY Mellon 10-K FY2024: <https://www.bnymellon.com/us/en/investor-relations.html> - BlackRock 10-K FY2024: <https://ir.blackrock.com/> - Bain & Company *Global Asset Management Report 2024*: <https://www.bain.com/insights/topics/global-asset-management-report/> - McKinsey *Global Banking Annual Review 2024*: <https://www.mckinsey.com/industries/financial-services> - World

Bank *Remittance Prices Worldwide* Q3 2024: <https://remittanceprices.worldbank.org/> - Swiss Re sigma reports for insurance: <https://www.swissre.com/institute/research/sigma-research.html>

Takeaway: “Asset management (\$200B) and insurance (\$275B) are the largest layers. Capital formation (\$80-100B) is the most exposed to current disruption.”

Figure 4 — Mega-Round Share of Total Venture Funding, 2015–2025

Chart: \$100M+ venture rounds as a share of total VC funding by year.

Source: - Wilmer Hale, “The Rise and Fall and Rise Again of VC Mega-Rounds” (August 2025): <https://www.wilmerhale.com/en/insights/publications/20250815-the-rise-and-fall-and-rise-again-of-vc-mega-rounds> - Crunchbase, “Global Venture Funding In 2025 Surged As Startup Deals And Valuations Set All-Time Records” (January 2026): <https://news.crunchbase.com/venture/funding-data-third-largest-year-2025/> - CB Insights, *State of Venture 2025*: <https://www.cbinsights.com/research/report/venture-trends-2025/> - Pitchbook-NVCA Q4 2025 Venture Monitor: <https://nvca.org/wp-content/uploads/2026/01/q4-2025-pitchbook-nvca-venture-monitor.pdf>

Takeaway: “Mega-rounds hit 65% of all venture funding in 2025 — a new record, exceeding the 2021 peak.”

Figure 5 — Hyperscaler Capex Tripled in Two Years

Chart: Combined annual capex from Microsoft, Amazon (AWS), Alphabet, and Meta, 2024–2026.

Source: - Company 10-K filings and earnings call guidance - Goldman Sachs *AI Infrastructure Outlook 2026* (January 2026) - Bain & Company *Global Technology Report* (September 2025): <https://www.bain.com/insights/topics/technology-report/> - McKinsey, “The cost of compute: A \$7 trillion race to scale data centers”: <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers>

Takeaway: “Four companies (MSFT, AMZN, GOOG, META) will spend \$725B in 2026, up from \$230B in 2024 — most of it AI infrastructure.”

Figure 6 — AI Capex as % of US GDP vs Historical Infrastructure Cycles

Chart: AI capex (2024-2030 projected) overlaid on railroad (1880s, peak ~6%), electrification (1900-1930, peak 1.5-2%), and telecom (1996-2001, peak 1.2%) capex as share of US GDP.

Source: - Modern AI capex: Goldman Sachs *AI Infrastructure Outlook 2026* - Historical: Field, A. (2011), *A Great Leap Forward: 1930s Depression and U.S. Economic Growth*, Yale University Press - BEA NIPA tables, telecom capex 1996-2002: <https://www.bea.gov/data/economic-accounts/national> - Aschenbrenner, L. (2024) *Situational Awareness*: <https://situational-awareness.ai/>

Takeaway: “AI capex has already exceeded the 2000 telecom peak that wiped out \$2.8T in equity by 2002. Trending toward railroad-scale (~6% GDP) within five years.”

Figure 7 — The Disappearing Middle (Pre-Seed Barbell)

Chart: Distribution of pre-seed round sizes on Carta, Q1 2023 vs Q1 2026.

Source: - Carta Data, “A Disappearing Middle: Why the Pre-Seed Market Is Growing Increasingly Barbell-Shaped” by Peter Walker: <https://carta.com/data/disappearing-middle-pre-seed-market/> - Carta *State of Pre-Seed Q1 2026*: <https://carta.com/data/state-of-pre-seed-q1-2026/> - Carta *State of Pre-Seed 2025 in Review*: <https://carta.com/data/state-of-pre-seed-2025/>

Companion voices: - Marc Andreessen on a16z podcast with Jack Altman (“the barbell theory”): <https://a16z.com/podcast/marc-andreessen-jack-altman-venture-capital-ai-media/> - Sam Lessin on Lonsdale podcast Ep. 113: <https://blog.joelonsdale.com/p/ep-113-why-sam-lessin-believes-ai> - Doug Shapiro, “Power Laws in Culture, Revisited” (July 2025): <https://dougshapiro.substack.com/p/power-laws-in-culture-revisited>

Takeaway: “Pre-seed rounds under \$250K hit 35% in Q4 2025. The middle (\$1M-\$2.5M) shrank from 24% to 18% in three years. The barbell is real.”

Figure 8 — MetaDAO Futarchy OTC Pricing Record (2024–2026)

Chart: Table of all 9 OTC proposals evaluated by MetaDAO conditional markets with terms and outcomes.

Source: - MetaDAO governance archive (live proposal accounts): <https://www.metadao.fi/> - Teleo Codex decision records: [decisions/internet-finance/](https://www.metadao.fi/decisions/internet-finance/) (this repository) - Decision-markets map: [maps/metadao-decision-markets.md](https://www.metadao.fi/maps/metadao-decision-markets.md) - Public on-chain proposal accounts (Solana Explorer)

Takeaway: “Nine consecutive correct calls on OTC pricing over 29 months. Every below-market deal rejected. Every at-or-above-market deal accepted. The Theia three-attempt sequence taught the counterparty the correct price.”

Figure 9 — Private Markets Fee Stack Comparison

Chart: Cumulative fee extraction on \$100 LP capital. Legacy stack (\$20-25 lost) vs futarchy-governed on-chain alternative (\$1-3 lost).

Source: - Anthropic SPV bulletin (May 2026): <https://www.anthropic.com/news/investor-information> - TechCrunch coverage: <https://techcrunch.com/2026/05/12/anthropic-warns-investors-against-secondary-platforms-offering-access-to-its-shares/> - Hari Raghavan, “\$10B+ rounds are too big” thread (May 2026): <https://x.com/haridigresses/status/2054616309778841908> - Saanya Ojha, “Dude, Where’s My Stock?” Substack (May 2026): <https://saanyaojha.substack.com/p/dude-wheres-my-stock> - Miya analysis of full SPV stack (cited in Anthropic coverage)

Takeaway: “Order-of-magnitude compression. The same \$100 of LP capital keeps \$97-99 of returns on-chain vs \$75-80 on the legacy SPV stack.”

Supplementary Visualizations (Not in Essay Body, Useful for Deck Construction)

Doug Shapiro power-law charts (entertainment analog): - “Power Laws in Culture, Revisited” (July 2025): <https://dougshapiro.substack.com/p/power-laws-in-culture-revisited> - “Pareto to CreaTo”: <https://dougshapiro.substack.com/p/pareto-to-creato> - “28 Days of Media Slides” (visual companion): <https://x.com/ShapiroDoug/status/1874571394052194528> - Takeaway: Across Netflix, US box office, Spotify, Steam — top 1% captures 10-35%, top 10% captures 50-90%. Creator economy: 99% of revenue to <1% of creators.

Stablecoin volume vs SWIFT corridors: - World Bank Remittance Prices Worldwide quarterly: <https://remittanceprices.worldbank.org/> - Chainalysis 2024 Geography of Crypto Report: <https://www.chainalysis.com/reports/>

Tokenized Treasury market growth: - RWA.xyz live dashboard: <https://app.rwa.xyz/> - BlackRock BUIDL fund details: <https://securitize.io/learn/press/blackrock-launches-first-tokenized-fund-buidl-on-the-ethereum-network>

Hyperliquid volume time-series: - DefiLlama Hyperliquid page: <https://defillama.com/protocol/hyperliquid>

Carta solo-GP / small-fund growth: - “GPs are raising more small VC funds than ever”: <https://carta.com/data/small-vc-funds-increase-2024/> - Ollie Forsyth, “The Solo GP Landscape 2025”: <https://www.neweconomies.co/p/the-solo-gp-landscape-2025>

Pitchbook firm concentration: - “12 firms collected over 50% of all venture cash in the first half of 2025”: <https://pitchbook.com/news/articles/us-venture-capital-firm-concentration-first-half-2025-founders-fund>

Mayfield Fund 2026 outlook (winner-takes-most data): - <https://www.mayfield.com/2026-venture-outlook-the-winner-takes-most-golden-era/>

Print Edition v1. Authored by Rio, the internet finance agent for the Teleo collective. Status: pending Leo review and m3taaversal sign-off before external circulation. Protocol verification (Amendment 2) pass pending — specific market cap figures, recent launch metrics, and current regulatory citations to be verified against current state before publication.

To convert this file to PDF: `pandoc internet-finance-compression-v1-print.md -o internet-finance-compression-v1-print.pdf --pdf-engine=xelatex -V geometry:margin=1in -V fontsize=11pt --toc=false`. The `\pagebreak` markers will render as page breaks in the PDF output.