

The Economics of Good Enough

Key Points

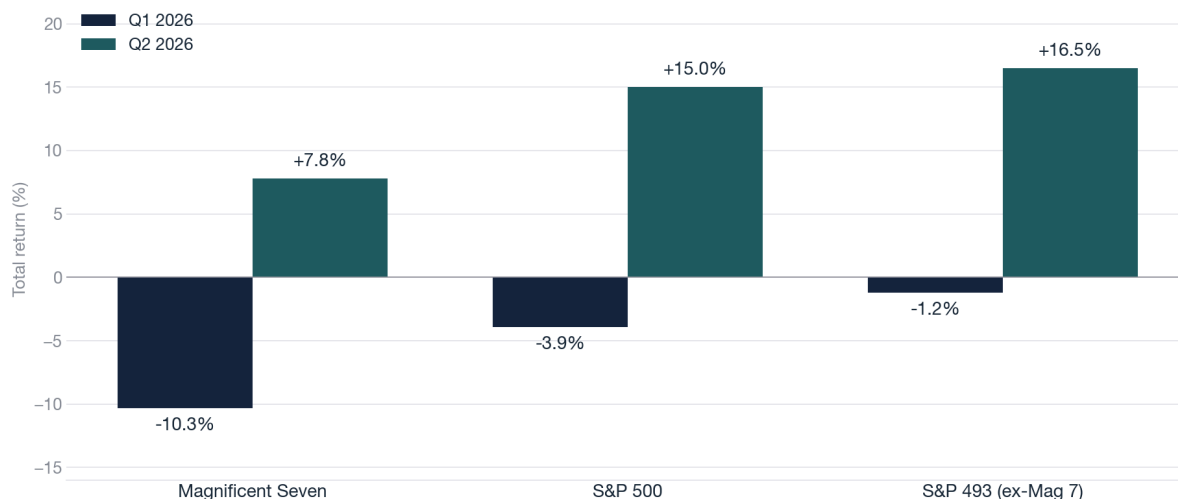
- 1. The S&P 500 rebounded roughly 15% in Q2 — its best quarter in six years — reversing the first-quarter selloff as the Iran conflict de-escalated and oil returned to pre-war levels.*
- 2. The rally was again led by artificial intelligence, but from beyond the mega-caps: memory and semiconductor names outside the Magnificent Seven drove the advance, while the Mag 7 themselves remain down for the year.*
- 3. Valuations are back near record extremes — a Shiller CAPE around 41, the second-highest reading in history — even as the Federal Reserve turned hawkish and pushed rate cuts off the table.*
- 4. The market spent the quarter celebrating the AI trade at precisely the moment the evidence turned against its economics: competition is competing away the pricing power the build-out assumes it will earn.*

Q2 2026 Market Performance

If the first quarter was defined by fear, the second was defined by relief. The S&P 500 rose roughly **15%** — its strongest quarter in six years — recovering nearly all of the first quarter's losses and closing the first half of 2026 modestly higher. The proximate cause was the unwinding of the shock that had driven the March decline: as back-channel negotiations gave way to a pause in hostilities between the United States, Israel, and Iran, and as commercial traffic resumed through the Strait of Hormuz, the geopolitical risk premium that had been hastily priced into equities in March was just as hastily priced back out.

The Round Trip

Total return by group — the Q1 2026 selloff and the Q2 2026 rebound



Source: Yahoo Finance; Cornell Capital Group calculations. Mag 7 market-cap weighted; S&P 493 = S&P 500 ex-Mag 7.



The reversal was most visible in energy and volatility. Brent crude, which had surged past \$100 and briefly touched \$126 during the March oil shock, fell back to roughly \$72 by late June — with West Texas Intermediate near \$69, levels last seen the day before the conflict began. Energy, the standout sector of the first quarter, became the laggard of the second. The VIX, which had spiked above 31 in March, settled back to around 16. In the space of three months the market had traveled from pricing a war to pricing its absence.

But the more important story lay beneath the headline recovery. The rebound was led, once again, by artificial intelligence — yet this time the leadership broadened well beyond the mega-caps. Memory and semiconductor companies, riding an AI-driven supply shortage, posted some of the largest quarterly gains on record: shares of one major memory producer more than tripled in the quarter. The Magnificent Seven participated but lagged, recovering only part of their first-quarter decline and finishing the first half still **down for the year**. The average stock, in other words, outpaced the index — a healthier internal dynamic than the narrow leadership of recent years, but one built almost entirely on a single thesis. Not every name shared in the optimism: several companies that cut guidance or announced restructurings were punished severely, a reminder that beneath a buoyant index, dispersion is widening.


The one clearly sober voice was the Federal Reserve. At its June meeting the FOMC held its target range at **3.50–3.75%** and, notably, shifted its guidance in a hawkish direction — with the elevated energy prices of the spring feeding into inflation expectations, the median official now sees no rate cut in 2026, and some see a hike. This is a meaningful change from the start of the year, when one or two cuts were the consensus. A market that rallied 15% did so into a policy backdrop that grew *less* accommodative, not more.

The Valuation Backdrop

The consequence of a 15% quarter layered on top of an already-expensive market is that valuations have returned to the extremes we flagged in our first-quarter letter. The Shiller cyclically-adjusted price-to-earnings ratio now sits around **41** — roughly double its long-run median of 16, higher than the 1929 peak, and second only to the March 2000 top in more than a century of data.

Second-Highest in History

Shiller cyclically-adjusted P/E (CAPE) at major U.S. market peaks vs. today



Source: Robert J. Shiller data, multpl.com; Cornell Capital Group. CAPE = price / 10-yr avg inflation-adjusted earnings.



We wrote last quarter that the equity risk premium — the compensation investors receive for bearing equity risk above the risk-free rate — had compressed to levels reminiscent of the late 1990s, and that the problem was not merely a *lower* premium but an *inadequate* one. Nothing this quarter improved that arithmetic; the rally made it worse. With the ten-year Treasury yielding roughly **4.4%** and the Fed signaling higher-for-longer, the hurdle that equities must clear to justify their valuations has risen even as prices have climbed to meet a thinner and thinner margin of safety.

Two features of this valuation compound the concern. The first is concentration: market-value creation has never been more concentrated in a handful of names, and the index's fate is now unusually hostage to a small number of AI-levered mega-caps. The second is the *source* of the earnings supporting today's multiples. As Bill Hester of Hussman Strategic Advisors documented this spring, the improved-looking forward multiples of 2026 owe less to cheaper prices than to analysts marking earnings estimates steadily *upward* — embedding record profit margins and a near-universal assumption that there will be no losers from AI. That assumption is precisely what this quarter's evidence calls into question.

The Economics of Good Enough

Which brings us to the theme of this letter, developed at greater length in our companion publication, *The Economics of Good Enough*.

Long-time readers will recognize the framework. As we wrote in our second-quarter 2021 letter, the hallmark of a *big market delusion* is that every firm in an evolving industry rises together even though they are direct competitors — each priced as a future winner, a fallacy of composition, because the parts cannot exceed the whole. We watched it in social media and in electric vehicles; in our mid-2023 letter we suggested AI might be the third act. The governing principle is simple: a new technology rewards its investors only if it earns returns on capital above the cost of capital, and that requires **barriers to entry** durable enough to keep competition from driving those returns back down.

Three years and several trillion dollars later, the question "what are the barriers?" is no longer theoretical. The five largest U.S. technology companies lifted combined AI capital spending from roughly \$227 billion in 2024 to \$391 billion in 2025; 2026 guidance has since climbed toward **\$700 billion**, and Goldman Sachs projects a cumulative **\$7.6 trillion** by 2031.

The Arms Race Accelerates

Combined AI capital expenditures of the five largest U.S. technology companies



Source: Company filings and guidance; Goldman Sachs projections (2027). 2026 reflects announced guidance.



Spending on this scale is rational only if it buys something durable, and the thesis underneath it is that scale itself is the moat. The past eighteen months have steadily undermined that idea. Chinese open-weight models — DeepSeek and Qwen among them — now handle the majority of commercial workloads at prices **7 to 25 times below** the U.S. frontier labs; a capability that commands a premium today is matched by a model costing a fraction as much within a year. The hardware tells the same story: the rental rate for the prior-generation flagship AI chip has collapsed from roughly \$8 an hour in 2024 to under \$1 today. A barrier that must be rebuilt every twelve months at rising cost is not a barrier but a treadmill — and it has a corollary in the accounts. Hyperscalers depreciate AI hardware


over five to six years, but the annual product cadence argues for an economic life closer to three; on that basis Research Affiliates estimates only about **one-third** of 2026 hyperscaler capex is genuine capital formation, the rest merely replacing gear that is already obsolete. Earnings resting on six-year depreciation are, to that extent, overstated.

On the demand side, the industry has just lived through the end of subsidized intelligence. Flat-rate subscriptions concealed AI's true cost through 2025; in the first half of 2026 the major providers moved almost simultaneously to usage-based pricing, and the response was immediate — firms that had never metered usage found they had burned through annual budgets in months and imposed hard per-employee caps. Citrini Research named the shift from "tokenmaxxing" to **"tokenpanic."** Once every prompt carried a price, procurement asked the question that ends premium pricing in every maturing market: *is the cheaper option good enough?* For summarization, classification, routine code, and internal search, the answer is increasingly yes. A Bain & Company survey of 951 executives at companies above \$100 million in revenue captured the other half of the problem: AI cost savings have fallen well short of expectations, yet most are raising budgets again — 44% funding the next wave from savings that **have not yet materialized**. Bain's summary is blunt: *"The technology worked. The value didn't arrive."*

Set against this demand picture, the supply commitments are extraordinary. Planned AI data-center capacity is on the order of 190 gigawatts — enough that, even on conventional utilization, it would require roughly **\$1.75 trillion in annual revenue by 2030** to earn its keep, against a combined run rate at the two leading labs of about \$60 billion today.

The Revenue Mountain

Annual AI revenue required by 2030 to justify planned data-center capacity, vs. today's run rate



Source: Sightline Climate planned-capacity data (190 GW, ~\$12.5M per critical MW per year); company disclosures.



And that demand is, to a striking degree, circular: the hyperscalers invest in the frontier labs, the labs spend the proceeds on hyperscaler compute, and the hyperscalers buy the chips to supply it, each leg citing the others as proof of demand. We need not adopt the most bearish reading to make the narrow point — the revenue mountain must be climbed by firms whose pricing power is being competed away *while they climb it*. The build-out math assumes premium monetization; the competitive dynamics assume the premium away. Both cannot be true.

Structural Risks

The theme above bears directly on the largest single risk to the index: the AI capital cycle. The mega-caps that dominate the S&P 500 are diverting record sums of owner earnings into assets with three-to-five-year economic lives, in pursuit of revenue the competitive evidence suggests will carry thinner margins than the businesses funding it. When every player must invest simultaneously to avoid falling behind, the likely outcome is mutual margin compression rather than durable advantage. History — from fiber optics in the 1990s back to the railroads — is consistent: transformative infrastructure often creates enormous social value while destroying much of the capital that builds it. Investors should watch depreciation schedules and disclosed useful-life assumptions as closely as they watch revenue guidance; that is where this story will surface in the financial statements first.

Alongside it sits the fiscal position, which continues to deteriorate. The CBO projects a **\$1.9 trillion** deficit in fiscal 2026, with federal debt around 101% of GDP and net interest payments approaching \$1 trillion annually — more than the entire defense budget. As we have argued before, persistent deficits have been a powerful tailwind for corporate earnings, but ever-rising debt brings higher interest costs and refinancing risk. With the Fed now signaling higher-for-longer rather than the cuts the market expected in January, the arithmetic of financing that debt — and of discounting equity cash flows against it — has become less forgiving. And the Iran truce that lifted markets in June remains fragile; the same oil-price sensitivity that faded this quarter can return quickly.

Investment Implications

None of this requires selling everything and moving to cash. It does have practical implications, and they extend the framework of our recent letters.

Separate "will AI work?" from "will AI investments pay?" These are routinely conflated and have opposite sensitivities to competition. Intense competition makes the technology better and cheaper *and* makes the returns on the capital behind it worse. The quarter's rally rewarded the first question; the durable returns depend on the second.

Interrogate reported earnings. Where AI hardware is depreciated over five or six years against a three-year economic life, current earnings are borrowing from future write-downs. The gap between accounting depreciation and economic obsolescence is the single largest soft spot in mega-cap income statements today.

Apply the pricing-power test to every AI claim. The question is not whether a company uses AI or even leads in it, but whether it could raise prices without losing the workload to an alternative that is 90% as good at 10% of the cost. Very few participants in the chain can currently answer yes. Those that can — sole-source physical suppliers, businesses with genuine switching costs, owners of distribution — are where durable value is most likely to accrue.

Revise return expectations, and let diversification do its work. With the CAPE near 41 and the risk premium thin, the penalty for concentration in expensive U.S. large-caps is severe, and high-quality bonds offer yields competitive with plausible equity returns at a fraction of the volatility. This does not mean abandoning equities; it means ensuring every dollar of equity exposure is earning its keep.

Conclusion

None of this is an argument that artificial intelligence will fail. It is an argument that the technology's success and its investors' returns are separate questions joined by a single variable: barriers to entry. Competition is the mechanism by which AI's benefits reach the broader economy — cheaper intelligence, embedded everywhere, priced like a utility. It is also the mechanism by which the excess returns assumed in today's valuations are competed down to the cost of capital.

That is the quiet irony of the second quarter. The market staged its best rally in six years to celebrate a technology whose economics, on the evidence now in hand, are hardening *against* excess returns at precisely the moment its capital commitments are largest. As always, our approach remains grounded in comparing price with value, and in demanding adequate compensation for the risks being borne. The cheaper intelligence becomes, the better for the world — and the steeper the mountain for those who paid for it to be expensive.