



The Efficiency Agonist

We consumed GPUs. We consumed megawatts of power. We consumed hundreds of billions of dollars in CapEx building data centers that look less like server farms and more like sovereign nation-states. If 2023 and 2024 were about the "builders" (the hyperscalers and the utility grid), 2025 was the year the narrative finally trickled down to the "users."

And it turns out, the killer app for AI isn't a chatbot that writes poetry. It is a metabolic agent.

In medicine, GLP-1 agonists work not just by suppressing appetite, but by fundamentally altering how the body processes energy. In 2025, we saw AI act as an "Efficiency Agonist" for the S&P 500. For the better part of the last decade, and certainly during the ZIRP (Zero Interest Rate Policy) era, corporate physiology was governed by a linear law: Revenue \approx Headcount. If you wanted to sell 20% more widgets, service 20% more clients, or write 20% more code, you generally needed to hire 20% more bodies.

That correlation has broken.

We are witnessing a "decoupling" of revenue growth from expense growth. While the "Mag 7" continued their CapEx arms race, the rest of the S&P 493 began quietly injecting AI into their operations to treat the chronic condition of administrative bloat. The result wasn't a sudden explosion in top-line revenue—it was a stubborn, structural expansion of operating leverage that has confounded the bears.

To understand why profit margins are defying gravity, we have to look past the "GenAI" buzzwords and look at the plumbing of specific sectors. This is the anatomy of the metabolic shift.

The Great Subsidy: Eating Your Own Cooking

One of the most popular bearish arguments of 2025 has been the "Return on Invested Capital" (ROIC) problem. Who is going to pay for all these chips?

Critics predicted that the massive depreciation schedules associated with building AI infrastructure would weigh heavily on Tech sector profitability. But the data tells a different story.

Look at the scoreboard for Q3 2025. Information Technology margins didn't contract; they expanded significantly, jumping to 28.0% from 25.1% a year prior. This defies the laws of heavy manufacturing. Usually, when you spend hundreds of billions on new factories (or data centers), your margins dip before they recover.

Why the divergence? Because the "Mag 7" are practicing what they preach. They are the primary users of their own efficiency tools. They are "eating their own cooking," using AI to automate their own code generation, customer support, and content moderation.

But the broader story is the deflationary subsidy this creates for the rest of corporate America. The Tech sector is building the railroads; the rest of the S&P 500 are just buying the tickets. Because the cost of intelligence (inference) has dropped exponentially, non-tech companies can now access "intelligence on tap" for pennies on the dollar.

This creates what macro strategists call the "CapEx-Opex Arbitrage." The massive capital expenditures of the Tech giants become the operating expense savings of everyone else. We see this clearly in the sector breakdown: Financials (margins up to 20.2% from 17.9%) and Utilities (up to 17.1% from 14.8%) are effectively leveraging this new infrastructure to widen their spread.

The "Junior Developer" Arbitrage (SaaS & Tech Services)

Nowhere is this shift more visible—and measurable—than in the software sector.

For twenty years, the "Junior Developer" has been a necessary loss leader. Fresh computer science graduates are expensive to recruit, slow to code, require heavy senior supervision, and are prone to bugs. They are an investment in the future, but a drag on the present.

In 2025, the "Copilot" seat license didn't replace the Senior Architect, but it largely automated the grunt work of the Junior.

Data from developer platforms this year has been staggering. Developers using AI assistants are completing tasks 55% faster than those without. But the real story isn't just speed; it's "Time-to-Merge"—the time it takes for a piece of code to be written, reviewed, and integrated into the main software product. That metric has improved by nearly 50% across the industry.

This has created a massive arbitrage opportunity. A Junior Developer costs a firm roughly \$150,000 per year (salary + benefits + overhead). A Copilot license costs roughly \$400 per year. Even if the AI is only 20% as effective as a human (and the data suggests it is much higher), the math is undeniable.

We are seeing enterprise software firms where maintenance coding, QA testing, and documentation—tasks that previously consumed 40% of engineering hours—are being offloaded to LLMs. This hasn't resulted in mass firings (yet); instead, it has resulted in a massive increase in Product Velocity per Dollar of SG&A.

Companies are effectively shipping twice as much product updates for the same payroll expense. The "R&D" line on the income statement is no longer just "Salary"; it is now "Salary + Compute," and Compute is getting cheaper while Salary is getting more expensive.

The "Klarna Effect": Support & Service

If Software is the "Junior Developer" story, Customer Service is the "Klarna" story.

The fintech payments sector became the "Patient Zero" for this thesis when early adopters revealed that their AI assistants were handling the workload of hundreds of full-time human agents.

The numbers were brutal for the "Human-in-the-Loop" proponents:

- Resolution Time: Dropped from 11 minutes to 2 minutes.
- Customer Satisfaction: Parity with human agents.
- Cost Savings: Millions in annualized savings.

In 2025, we saw this "Klarna Effect" spread to the S&P 500 industrials and carriers. It wasn't just about chatbots on a website; it was about internal support.

Consider the P&C Insurance sector. Historically, this business is a game of friction. A spike in claims (due to a hurricane or freeze) usually necessitated a spike in adjusters and temporary support staff, eroding the Combined Ratio.

This year, major carriers deployed multi-modal models capable of "triaging" claims. These systems analyze photos of fender benders or water damage, cross-reference policy limits, and estimate payouts with 95% accuracy in seconds. The human adjuster now only touches the complex anomalies—the "edge cases."

Analytics firms report that routine claims processing times dropped from 7-10 days to 24-48 hours. For an insurance carrier, time is money—literally. The faster a claim is closed, the lower the administrative burden and the higher the customer retention.

The result? A carrier can now grow its premiums written by 15% without adding a single square foot of back-office call center space. That is pure margin expansion.

The End of the "Billable Hour" Bloat (Professional Services)

Perhaps the most culturally significant shift is happening in the legal and consulting worlds—the bastions of the "Billable Hour."

For decades, law firms and consultancies thrived on the inefficiency of discovery and document review. It was a volume business. If a client was sued, you threw 50 associates at the problem to read emails for six months. You billed for every hour.

In 2025, the clients woke up.

Corporate legal departments, armed with their own internal LLMs, started pushing back. Why pay \$800 an hour for a first-year associate to summarize a deposition when an AI can do it in 30 seconds?

Industry reports suggest that AI tools can save approximately 190 work-hours per lawyer per year. Across the US legal market alone, that represents a \$20 billion efficiency dividend.

While this threatens the revenue model of the law firms (who are frantically pivoting to "value-based billing"), for the clients -the S&P 500 corporations- it is a massive deflationary force on their legal and compliance lines. What used to be a \$5 million external counsel spend for a merger review is becoming a \$2 million spend.

That \$3 million difference drops straight to the bottom line of the acquirer. Multiply that across thousands of corporate transactions, audits, and compliance checks, and you begin to see why SG&A expenses are trending down as a percentage of revenue.

The "Sticky Margin" Thesis

This phenomenon explains one of the year's biggest conundrums: Why haven't profit margins mean-reverted?

Classic economic theory screams that record-high profit margins act as a magnet for competition. If margins are high, new entrants come in, undercut prices, and drive margins back to the historical average. Bears have been betting on this reversion since 2022.

But they missed the productivity shock.

Quantitative research now suggests that the old rules of mean reversion are broken because the composition of the index has changed. We are no longer an index of steel mills and grocery stores; we are an index of asset-light, high-leverage monopolies.

But even beyond composition, AI is acting as a margin defense mechanism.

If AI allows a company to maintain pricing power while structurally lowering its cost to serve (the "Efficiency Agonist" effect), margins don't revert—they reset higher. The "efficiency gains" we promised in the pitch decks of 2023 finally showed up in the 10-Ks of 2025.

When a company can increase output without increasing headcount, its operating leverage goes parabolic. We are seeing early signs of this in the "Revenue per Employee" metrics of the S&P 500, which have ticked up significantly after stagnating in the early 2020s.

Conclusion: Leaner, Not Just Larger

As we evaluate equity valuations for 2026, the metric of choice is shifting. We have stopped asking just "How much will they grow?" and started asking "How lean can they get?"

The companies winning in this environment aren't just the ones selling the chips; they are the ones taking the medicine. They are using this technology to sever the link between growth and operational complexity.

The risks, of course, are human. The "Efficiency Agonist" implies a reduction in intake, and in corporate speak, that means fewer jobs (or at least, fewer new jobs). This creates a macroeconomic tension: Can consumer spending hold up if the corporate sector stops hiring?

So far, the answer has been yes, largely because the demographics of the labor market (Baby Boomer retirements) have created a natural shortage of workers. AI isn't "taking jobs" so much as it is filling the hole left by a shrinking workforce. It is the only way we can grow GDP in a world with flat demographics.

But for the shareholder, the conclusion is clear. The "AI Trade" has moved from the construction phase to the harvest phase. The builders have built. Now, the users are getting slim. And a slimmer S&P 500 is a much faster runner.

IMPORTANT LEGAL DISCLOSURES

CURRENT MARKET DATA IS AS OF 12/30//2025. OPINIONS AND PREDICTIONS ARE AS OF 12/30/2025 AND ARE SUBJECT TO CHANGE AT ANY TIME BASED ON MARKET AND OTHER CONDITIONS. NO PREDICTIONS OR FORECASTS CAN BE GUARANTEED. INFORMATION CONTAINED HEREIN HAS BEEN OBTAINED FROM SOURCES BELIEVED TO BE RELIABLE BUT IS NOT GUARANTEED.

THIS PRESENTATION (THE "PRESENTATION") HAS BEEN PREPARED SOLELY FOR INFORMATION PURPOSES AND IS NOT INTENDED TO BE AN OFFER OR SOLICITATION AND IS BEING FURNISHED SOLELY FOR USE BY CLIENTS AND PROSPECTIVE CLIENTS IN CONSIDERING GFG CAPITAL, LLC ("GFG CAPITAL" OR THE "COMPANY") AS THEIR INVESTMENT ADVISER. DO NOT USE THE FOREGOING AS THE SOLE BASIS OF INVESTMENT DECISIONS. ALL SOURCES DEEMED RELIABLE HOWEVER GFG CAPITAL ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES. THE OPINIONS CONTAINED HEREIN ARE NOT RECOMMENDATIONS.

THIS MATERIAL DOES NOT CONSTITUTE A RECOMMENDATION TO BUY OR SELL ANY SPECIFIC SECURITY, PAST PERFORMANCE IS NOT INDICATIVE OF FUTURE RESULTS. INVESTING INVOLVES RISK, INCLUDING THE POSSIBLE LOSS OF A PRINCIPAL INVESTMENT.

INDEX PERFORMANCE IS PRESENTED FOR ILLUSTRATIVE PURPOSES ONLY. DIRECT INVESTMENT CANNOT BE MADE INTO AN INDEX. INVESTMENT IN EQUITIES INVOLVES MORE RISK THAN OTHER SECURITIES AND MAY HAVE THE POTENTIAL FOR HIGHER RETURNS AND GREATER LOSSES. BONDS HAVE INTEREST RATE RISK AND CREDIT RISK. AS INTEREST RATES RISE, EXISTING BOND PRICES FALL AND CAN CAUSE THE VALUE OF AN INVESTMENT TO DECLINE. CHANGES IN INTEREST RATES GENERALLY HAVE A GREATER EFFECT ON BONDS WITH LONGER MATURITIES THAN ON THOSE WITH SHORTER MATURITIES. CREDIT RISK REFERES TO THE POSSIBLITY THAT THE ISSUER OF THE BOND WILL NOT BE ABLE TO MAKE PRINCIPAL AND/OR INTEREST PAYMENTS.

THE INFORMATION CONTAINED HEREIN HAS BEEN PREPARED TO ASSIST INTERESTED PARTIES IN MAKING THEIR OWN EVALUATION OF GFG CAPITAL AND DOES NOT PURPORT TO CONTAIN ALL OF THE INFORMATION THAT A PROSPECTIVE CLIENT MAY DESIRE. IN ALL CASES, INTERESTED PARTIES SHOULD CONDUCT THEIR OWN INVESTIGATION AND ANALYSIS OF GFG CAPITAL AND THE DATA SET FORTH IN THIS PRESENTATION. FOR A FULL DESCRIPTION OF GFG CAPITAL'S ADVISORY SERVICES AND FEES, PLEASE REFER TO OUR FORM ADV PART 2 DISCLOSURE BROCHURE AVAILABLE BY REQUEST OR AT THE FOLLOWING WEBSITE: [HTTP://WWW.ADVISERINFO.SEC.GOV/](http://www.adviserinfo.sec.gov/).

ALL COMMUNICATIONS, INQUIRIES AND REQUESTS FOR INFORMATION RELATING TO THIS PRESENTATION SHOULD BE ADDRESSED TO GFG CAPITAL AT 305-810-6500.