

VALUE INVESTING AND THE EMERGING 4.0

By: Christopher Tsai

May 2026

Technological advancements have dramatically transformed growth, innovation, and value creation over the past few decades. Traditional investing mindsets are lagging behind this shift. As David Foster Wallace reminded us in his “This is Water” parable, the most obvious realities are often the hardest to see—the invisible “water” of assumptions and dogmas in which we swim. Investors have been evolving their approach to value investing for decades, but it is time to take the next step: from hunting bargains in dying assets to backing the visionaries shaping tomorrow. Here’s how value investing has evolved through 1.0, 2.0, 3.0, and the emerging 4.0.

Value Investing 1.0: Emerging around the Great Depression, Value Investing 1.0 was all about buying assets at a steep discount to book value, or even net current asset value. Think of the latter as buying the entire company, factory and all, for less than two-thirds of its net liquid assets alone—getting the plant, property, and equipment for free. Benjamin Graham outlined this approach in the 1934 edition of *Security Analysis*. This era focused intensely on current tangible metrics. Warren Buffett later nicknamed these deeply discounted opportunities “cigar butts.”

Value here was immediate: it resided entirely in the company’s then-current balance sheet. Recent earning power was a bonus. This deep value approach was very much a product of its time: a manufacturing-heavy world filled with tangible assets that depreciated slowly, with plenty of forced liquidations creating genuine bargains. It worked brilliantly back then but downplayed growth potential.

Most investors still fixate on a single year of earnings—whether the past twelve months or the forward twelve months—and judge a company’s valuation based on those figures alone, rather than thinking in terms of “tomorrow to book.” In today’s hyper-connected economy, this approach misses how the most transformative opportunities lie not in comforting low price-to-earnings ratios or discounts to today’s book value, but in forward-thinking domains where real multi-baggers are born.

Value Investing 2.0: Recognizing these limitations, and needing to deploy increasing amounts of capital, Warren Buffett and Charlie Munger pioneered the shift from “cheap” to “quality at a fair price.” Buffett crystallized the new philosophy when he declared it is “far better to buy a wonderful company at a fair price than a fair company at a wonderful price.” The focus turned to moats—durable competitive advantages like powerful worldwide brands, network effects, or scale—that protect earnings and allow companies to compound at above-average rates over the long term.

Value here became a function of the present value of discounted cash flows, extending meaningfully into the future and farther out than the immediate focus of 1.0. With Value Investing 2.0, a margin of safety at purchase remains essential, committing capital only when the market offers a large discount to intrinsic value, calculated primarily through discounted cash flow analysis projecting future free cash flows. This mindset remains relevant in many industries, but it can falter amid relentless disruption, where near-term earnings metrics fail to capture the full picture.

Value Investing 3.0: Value Investing 3.0 offers a powerful lens for investing in growth companies. It builds on 2.0 but fully adapts to platforms where reinvestment is deliberately prioritized over short-term profits.

Drawing from Marshall McLuhan, the new “medium” (the environments and ecosystems engineered by technology) becomes the ultimate moat: “the medium is the moat.” Leading companies don’t merely sell products or services; they architect entire digital infrastructures and economic castles that redefine how we live, work, and interact—often invisibly. These moats form through self-reinforcing network effects and data dominance that compound value durably, following fractal-like, non-linear patterns of scaling that Benoit Mandelbrot’s work on fractals and power-law distributions helps us appreciate. Extreme inequality in outcomes is the rule, not the exception: a few dominant players capture outsized rewards as growth accelerates at higher scales.

I seek companies with bold visions operating in massive markets. They channel substantial investments into R&D, customer acquisition, operations, and ecosystems today—often temporarily depressing current earnings through heavy yet efficient reinvestment—precisely to create even greater shareholder value tomorrow. Cash flows are pushed out even further, with much of the intrinsic value now residing in the distant terminal value.

A Particularly Powerful Subset: Supply-Side and Demand-Side Economies of Scale

While many compelling 3.0 businesses operate through various forms of ecosystem building, a select and especially potent subset harnesses scale-economies-shared business models. These models represent one of the most dynamic expressions of the 3.0 era. I am grateful to Nick Sleep for first bringing the powerful concept of scale-economies-shared to the investing public; his insights on businesses that deliberately reinvest scale advantages back into customers rather than extracting maximum short-term margins have been enormously influential. Building on that foundation, I believe the picture is even richer when one distinguishes between its two fundamental expressions.

On the supply side, companies produce goods or services that carry a positive marginal cost. As they scale, unit costs fall dramatically and profitability rises. Rather than pocketing every incremental margin, these companies deliberately pass a substantial share of the benefit back to consumers in the form of lower prices. Amazon, Tesla, and SpaceX are vivid examples.

Amazon generously redistributes margins to consumers, igniting a virtuous flywheel that perpetually strengthens its ecosystem. Tesla strategically reduces vehicle prices to accelerate adoption, broaden its total addressable market, and gather ever more real-world data. Starlink has strategically lowered hardware costs and monthly service fees while dramatically improving speeds and network capacity, delivering far more data and bandwidth per dollar and accelerating global broadband adoption.

In every case, today’s strategic generosity creates an easily visible positive feedback loop: lower prices drive greater adoption and scale, which in turn funds even lower costs and stronger moats.

Yet supply-side economies are only half the story. Many of the most powerful platforms in the software and AI era operate with zero (or near-zero) marginal cost once the infrastructure is built. When the price of a service is already zero to the end user, it cannot be reduced further. In this environment, the company instead delivers progressively greater value through a dramatically improved user experience. I have coined the term “Engagement Value per Unit Time” to capture this demand-side dynamic. This enhanced value can take many forms: greater personalization and intelligent recommendations, heightened performance and reliability, new formats and features, and a generally more intuitive and user-friendly experience.

These demand-side economies create equally powerful—though subtler—positive feedback loops. Engagement Value per Unit Time draws more users to the platform and encourages each user to spend more time there. Those users, in turn, contribute their own data, creativity, and interactions, which further enrich the platform and improve the experience for everyone else. The result is a self-reinforcing flywheel that is harder to discern than the obvious price reductions of supply-side models, yet no less potent. Platforms such as YouTube, Instagram, and Google Search compete fiercely on this dimension, not on price.

In both supply-side and demand-side economies of scale, latent pricing power is created. This pricing power can be exercised—or flexed—to varying degrees (a little or a lot) as the business matures. The extent to which a company chooses to flex its latent pricing power has profound implications for valuation. What might appear expensive on the surface often is not, precisely because the company has intentionally refrained from exercising that power in order to invest more deeply in the customer value proposition. Across all scale-economies-shared businesses, this deliberate investment in the customer extends the runway for durable, compounding growth and builds economic castles that become ever more resistant to competition.

A Note on Survivorship Bias

The examples above—Amazon, Tesla, Google—are the survivors. For every company that successfully executed a 3.0 strategy, many more attempted it and failed: they misread network-effect dynamics, over-invested in the wrong ecosystem, or ran out of patient capital before reaching escape velocity. The 3.0 lens is powerful precisely because it is difficult to apply correctly. Identifying genuine ecosystem architects—as opposed to merely capital-intensive businesses with weak moats—demands rigorous due diligence, channel checks, and a sober assessment of competitive positioning across multiple scenarios. The framework is a starting point, not a shortcut.

Value Investing 4.0: This emerging era builds directly on 3.0 but pushes the time horizon of value creation even further. While a 3.0 company already has an operating business (the “node” or “nodes”) running atop an established foundation (the “edge”), a 4.0 company must first build the edge itself—the foundational infrastructure—before the operating company (the “node”) can even exist. When the edge and nodes come together, they form a powerful graph. Because these efforts require enormous capital and long periods of losses, 4.0 businesses are more likely to be incubated as subsidiaries or divisions within already-profitable 3.0 companies that can fund the foundational bets.

According to the definition developed in this framework, a true 4.0 business has:

- Intelligence itself as the economic castle (the “edge”)
- Digital labor as the core output consumers buy
- A self-sustaining model
- The ability to invest enormous capital into operating businesses (the “nodes”) at high ROIC

True 4.0 businesses do not yet exist in mature form. I am not aware of any independent, publicly traded (or even large private) company whose primary economic engine is a self-sustaining, autonomous intelligence system that generates digital labor as its core output.

The closest examples of 4.0 businesses are all incubated subsidiaries or heavily backed divisions of larger 3.0 companies.

Examples of 4.0 Businesses

Effort	Type	Status	Capital Source
Tesla (Dojo + FSD + Optimus)	Incubated subsidiary	Most advanced public example	Funded by profitable car and energy businesses
Frontier AI labs (Anthropic, OpenAI, xAI)	Heavily backed divisions	Private, heavy R&D phase	Funded by large 3.0 technology companies

Early signals of this emerging era are visible in these pioneering efforts. In Value Investing 4.0, the economic castle is intelligence itself. And for the first time, the primary output that consumers are buying is digital labor—autonomous intelligence systems performing work, making decisions, and delivering value on their behalf.

Companies are not simply building ecosystems or platforms—they are architecting self-improving, autonomous intelligence systems that create powerful positive feedback loops of recursive data, models, compute, and real-world feedback. These loops, further amplified by Metcalfe’s Law, enable the intelligence to compound at machine speed and scale.¹

These businesses are likely to generate little to no reported profit for years while they pour capital into foundational intelligence capabilities. The payoff, when it arrives, can be enormous—akin to a biotech company that invests heavily in R&D with no earnings for a decade, only to deliver a breakthrough therapy that creates generational wealth. The difference in 4.0 is that the “therapy” is intelligence itself: fleets of AI agents that allocate capital, serve customers, discover new products, and even redesign their own architectures.

Valuation requires an even more probabilistic, optionality-rich mindset—projecting not merely five- or ten-year cash flows, but the present value of capabilities whose impact may reshape entire industries in ways we cannot yet fully model.

For example, an investor might treat the historical base-rate success as a prior probability, then apply Bayesian updating to revise the posterior probability distribution as new evidence arrives—such as model performance, revenue traction, deployment milestones, or competitive developments. This produces a continuously refined range of possible terminal values rather than a single point estimate. The risk, of course, is profound uncertainty: we cannot know with precision what the future IRR on that incremental capital will ultimately be, demanding disciplined probabilistic valuation and a strict margin of safety particularly in highly uncertain environments.

¹ Metcalfe’s Law posits that the value of a network grows proportionally to the square of the number of its connected users or nodes, capturing the exponential power of network effects. In the context of companies like Tesla and SpaceX, this principle highlights how interconnected ecosystems—such as Tesla’s data-sharing vehicle fleet or Starlink’s satellite constellation—amplify value as more participants join, creating self-reinforcing moats that enhance scalability, resilience, and competitive dominance, often remaining invisible to participants (like water to a fish).

History offers a sobering guide: the base-rate success of companies pursuing each successive era’s strategy shrinks sharply—from roughly 60–70% for disciplined 1.0 practitioners to perhaps 1–5% for 4.0 pioneers—which argues not against investing in frontier intelligence, but for explicit scenario-weighting, position sizing commensurate with the probability of loss, and the intellectual honesty to distinguish a genuinely differentiated architecture from a well-funded experiment.

Progression of Cash-Flow Time Horizon Across Eras

Era	Time Horizon	Primary Valuation Focus	Key Mindset Shift	Est. Base-Rate Success*
1.0	Very immediate	Current balance sheet + recent earnings	Buy statistically cheap assets	~60–70%
2.0	Near-to-medium term (DCF)	Quality moats at a fair price	Buy wonderful businesses	~30–40%
3.0	Long-term, heavy terminal value	Future ecosystem power & compounding	Reinvest aggressively today	~5–15%
4.0	Extremely long-term / highly uncertain	Probabilistic valuation of autonomous intelligence & digital labor	Invest in the 'edge'	~1–5%

* Estimated base-rate success refers to the proportion of companies (or strategies) pursuing each era’s approach that achieved durable, compounding shareholder returns for a decade or more. These are rough order-of-magnitude estimates informed by: (1) historical backtests of Graham-style net-net investing (e.g., Carlisle, Mohanty & Oxman, 2010); (2) long-term active fund performance data (SPIVA/Morningstar reports showing the majority of active managers underperform over 10+ years); and (3) venture capital and corporate longevity studies (e.g., ~75% of VC-backed companies never return capital; only ~9% of VC companies drive 100% of the gains—Adams Street Partners; see also Bessembinder’s work on the extreme concentration of stock returns and corporate longevity). The numbers are intended to calibrate ambition rather than serve as precise forecasts. Source: Tsai Capital

Conclusion

I still like Value Investing 2.0 a great deal, while keeping a watchful eye on the emerging 4.0 frontier. That said, both 3.0 and especially 4.0 carry added risks. Many such companies become overvalued when hype overshadows substance. Spotting true ecosystem builders—and the even rarer intelligence architects—demands rigorous due diligence. Yet the most transformative firms often hide immense value beneath the surface, ignored by investors still fixated on near-term P/E ratios that cannot see the invisible water in which we all swim.

The base-rate success rates in the progression table are a necessary corrective to optimism: as the time horizon of value creation extends, the proportion of attempts that succeed shrinks sharply. This does not argue against investing in 3.0 or 4.0 opportunities—it argues for selectivity, position sizing commensurate

with the probability of loss, and the intellectual honesty to distinguish a genuinely differentiated intelligence architecture from a well-funded experiment.

The lesson: In exponential tech, value lives in future earnings power—often years out, hidden in foundational bets that depress today's profits. Vision, patience, and the ability to see the invisible "water" around us can turn what appears "overvalued" into generational wealth, but only when paired with rigorous probabilistic analysis, an explicit margin of safety, disciplined scenario-weighting, and the humility to acknowledge what we do not yet know.

Important Disclosures

Past performance is no indication or guarantee of future performance and no representation or guarantee is being made as to the future investment performance of Tsai Capital's separately managed accounts or any entity.

The information provided herein includes certain opinions, portfolio and/or market characteristics, and other information provided by Tsai Capital, as of May 2026. Any opinions or other information included herein may change without notice and should not be relied upon in selecting any investment or investment-related product or service. Nothing contained herein should be construed as investment advice.

References herein to Tsai Capital's efforts to minimize losses and seek a margin of safety should not be construed to imply an absence of risk in any investment. All investments carry risk, including the risk of loss of investment principal. Additionally, short-term market volatility may present increased risks for investors who have shorter investment horizons due to impending or current liquidity needs.