BEYOND THE PERPETUITY TRAP

By: Christopher Tsai

July 2026

The discounted cash flow (DCF) model is one of the most powerful instruments in the analyst's kit—rigorous, theoretically grounded, and intellectually honest about what it requires. I have been thinking a great deal lately about a problem that I suspect most investors have encountered but rarely examined head-on: the mismatch between the tool we use and the question we are actually trying to answer. When the question is where a stock will trade in, say, five years, the DCF is the wrong instrument, pressed into service because it is familiar, not because it fits.

As David Foster Wallace observed, we can live inside certain assumptions so completely that we stop noticing they are assumptions at all.[1] In this paper, I argue that price targets require a different framework and explain why forward earnings multiples provide the answer. The consequences of mismatching our tools to our desired solutions are systematic and largely invisible—invisible in the way that the most obvious assumptions often are: not because they are hidden, but because they surround us.

I. The Shifting Time Horizon of Value

Value investing has always adapted to the dominant mode of wealth creation in its era. Benjamin Graham's framework emphasized tangible discounts to book value and net current assets—a precise calibration to a capital-intensive industrial economy where hard assets were measurable and depreciated slowly. Warren Buffett and Charlie Munger then pioneered the shift toward long-run earning power and the compounding economics of businesses with genuine competitive advantages. A third era, Value Investing 3.0, has grappled with businesses that deliberately suppress near-term earnings to build self-reinforcing platforms—network-effect businesses, data flywheels, and marketplace ecosystems whose value is often latent. And the most forward-looking investment work now confronts an even longer and more uncertain horizon: foundational capabilities in machine intelligence that must be built years before profitable and self-sustaining operating businesses can fully emerge.

I have written about this progression at some length in my Value Investing 4.0 framework,[2] where I argue that the unit of competitive advantage has migrated from physical assets to brands to network-reinforced platforms to foundational intelligence—and that each migration demands adapted tools for estimation and judgment. A framework calibrated for one era will often systematically misprice the dominant opportunities of the next. This is not a minor inconvenience; it is the central analytical challenge of our time. Yet through all of these shifts, one principle has remained constant in the practice of value investing:

[1] David Foster Wallace, "This Is Water" (commencement address, Kenyon College, Gambier, OH, May 21, 2005). Published posthumously as *This Is Water: Some Thoughts, Delivered on a Significant Occasion, about Living a Compassionate Life* (New York: Little, Brown and Company, 2009).

[2] Christopher Tsai, "Value Investing and the Emerging 4.0," *Tsai Capital*, May 2026, https://tsaicapital.com/value_investing_and_the_emerging_4.php

the necessity of a margin of safety. While the nature of competitive advantage and the tools required to evaluate it have evolved, the discipline of demanding a sufficient gap between price and value has not.

What makes this especially striking is that precisely as value-creation horizons have lengthened, the average lifespan of large public companies has continued to contract. Research from Innosight shows that the average tenure of companies in the S&P 500 has fallen into the 15-to-20-year range in recent years, down from 30 to 35 years in the late 1970s.[3] At current rates of churn, a substantial share of today's leading companies are expected to disappear within the next decade or two through bankruptcy, acquisition, or various forms of creative destruction. The forces driving this compression—technological S-curves, winner-take-most dynamics, and rapid platform substitution—are the same forces making future cash flows both more valuable in theory and more uncertain in practice. This is the precise environment in which any valuation framework built on perpetuity assumptions deserves the most scrutiny.

II. The Architecture and Fragility of DCF

A standard two-stage DCF projects free cash flow (FCF) explicitly for a finite period and then estimates the terminal value (TV) of all remaining cash flows, most commonly using the Gordon Growth Model.

$\mathbf{TV_n = FCF_{n+1} \div (r - g)}$ where r = required return on capital (typically WACC) and g = assumed perpetual nominal growth rate.

The model is elegant and internally consistent. But it is also a highly leveraged function of two numbers that no analyst can estimate with genuine confidence over a multi-decade horizon. Both the discount rate and the perpetual growth rate are highly subjective inputs, and small changes in either can produce large swings in value. When these two variables move together, the effect is even more pronounced—what Charlie Munger might have called a Lollapalooza.[4]

Consider a concrete illustration. Suppose a business generates $1 billion of free cash flow in year ten, growing at a long-run rate of 3 percent and discounted at 9 percent. The TV is $1 billion divided by 6 percent, or roughly $16.7 billion. If the assumed perpetual growth rate rises by a single percentage point to 4 percent, the TV increases to $20 billion—a 20 percent increase from a 100-basis-point change in a number that is, in practice, a guess about nominal growth rates a decade or more into the future.

A similar effect occurs when the discount rate changes. Lowering the discount rate from 9 percent to 8 percent while holding growth at 3 percent produces the same 20 percent increase in TV. In other words, a one-percentage-point change in the assumed weighted average cost of capital (WACC) has the same impact on TV as a one-percentage-point change in the perpetual growth rate. Because the TV in this example represents approximately 70 percent of total enterprise value, either shift moves the overall valuation by roughly 14 percent.

The model that appears to produce a precise intrinsic value is, beneath its precision, a sensitivity machine. DCF does not eliminate uncertainty. It packages it—invisibly—mostly into the TV, where it compounds quietly and rarely gets examined.

---

[3] S. Patrick Viguerie, Ned Calder, and Brian Hindo, *2021 Corporate Longevity Forecast: Creative Destruction Is Accelerating* (Innosight, May 2021); see also Apollo Academy, "Average Tenure of Companies in the S&P 500 Index: 15 Years," March 2026.

[4] Charlie Munger introduced the term "Lollapalooza effect" to describe the confluence of multiple forces reinforcing one another to produce an outsized outcome. See Charles T. Munger, *Poor Charlie's Almanack*, ed. Peter D. Kaufman (Virginia Beach: Donning Company Publishers, 2005).

This sensitivity is amplified by the fact that there is no single, objective WACC. The discount rate itself is highly subjective and varies significantly depending on how it is calculated. Different assumptions about beta, the equity risk premium, the cost of debt, the company's target capital structure, and even the risk-free rate can produce a materially different WACC for the same business.

NVIDIA Corporation offers a clear illustration. As of mid-2026, estimates of its WACC range from 9.72% to 15.94%.[5] Holding all other assumptions constant, these two rates can produce intrinsic values that differ by approximately 50 percent. Two thoughtful analysts can therefore reach quite different conclusions about the same company's worth, simply by applying different—but entirely defensible—discount rates. The choice of WACC is not a neutral technical input but a judgment call that carries substantial weight in the final result.

III. Price Targets Require a Different Lens

A DCF is designed to answer: what is this business worth today, assuming perpetual existence? But that is not the question a price target is trying to answer. Converting a DCF present value into a price target requires additional assumptions about how the market's required return, competitive dynamics, and investor sentiment will evolve between now and then.

More precisely, the compound return an investor earns if a stock reaches a price target over any finite horizon is determined by the gap between today's entry price and the target—it is not, in general, equal to the discount rate embedded in the model. The discount rate asks: what return justifies owning this asset forever? The compound return over a finite horizon asks: what will I earn over this specific window? These are different questions. Treating them as equivalent is a category error that produces systematic confusion in both target-setting and portfolio allocation.

A price target is a forecast of where the market will clear at a particular point in the future, given projected fundamentals and the multiple investors are then willing to pay. It is not a present-value calculation rolled forward. Markets discover prices through multiples far more often than through explicit perpetuity discounting—and this is observable in practice. A review of sell-side equity research suggests that the large majority of price targets are derived primarily from forward earnings or cash flow multiples, with DCF analysis typically serving as a secondary check rather than the primary instrument.[6]

While a sophisticated observer might note that forward multiples are, in effect, compressed forms of DCF analysis—embedding assumptions about growth and required returns no less than a terminal value does—this observation, while technically correct, misses the practical distinction. The advantage of multiples lies not in escaping perpetuity assumptions, but in making those assumptions more transparent, testable against market precedent, and open to challenge. By contrast, the critical assumptions in a DCF—particularly those embedded in the TV—are often opaque, difficult to interrogate, and rarely subjected to the same level of external scrutiny.

---

[5] Alpha Spread; ValueInvesting.io; Finbox; GuruFocus; and various analyst models, as of June 2026.

[6] This observation is consistent with both academic studies and practitioner surveys of sell-side equity research. See Paul Asquith, Michael B. Mikhail, and Andrea S. Au, "Information Content of Equity Analyst Reports," *Journal of Financial Economics* 75, no. 2 (2005): 245–82, which found that 99 percent of reports by *Institutional Investor* All-American analysts referenced earnings multiples while only 13 percent referenced discounted cash flow analysis. A more recent CFA Institute survey of nearly 2,000 equity analysts similarly found that 93 percent reported using a market multiples approach (cited in Morgan Stanley, *Valuation Multiples*, April 2024).

When serious investors set price targets, they typically project earnings or cash flow to the target date and apply a multiple that reflects their assessment of the business at that future point. DCF analysis can inform that judgment, but it does not replace it.

The challenge lies in selecting the multiple. A disciplined approach begins by being explicit about the characteristics the business is expected to exhibit at the target date: its growth rate, returns on invested capital, margin structure, competitive position, capital intensity, and risk profile.

With these attributes in view, the analyst can then look to two useful references. The first is the multiples at which comparable companies have traded when they displayed similar growth rates, returns on capital, and competitive dynamics. The second is the multiples at which the company under analysis has itself traded historically, particularly during periods when its growth, returns, and business characteristics were comparable to those projected in the future. The analyst should also take into account the prevailing interest rate environment, as it can materially influence the level of multiples that markets are willing to pay.

These two references rarely point in exactly the same direction, and reconciling them requires judgment. When the company's historical multiples are materially higher than those of current comparables, the analyst must ask whether the business has become structurally better—or whether the market is simply applying a higher multiple to a temporarily favorable environment. Conversely, when comparables trade at higher multiples than the company's historical multiples would suggest, the analyst should consider whether the market is correctly anticipating improvement or simply paying up for narrative momentum. In both cases, the goal is not to average the two data points mechanically, but to understand the underlying drivers of any divergence and to form a view on which reference is more relevant to the future state being projected.

While neither reference is perfect, the analyst can synthesize them to establish a reasonable range for the exit multiple, adjusting for differences in quality, durability of competitive advantage, and prevailing market conditions. Together, these references provide a more grounded basis for the exit multiple than speculation or simple extrapolation.

A target built from forward earnings and a well-reasoned exit multiple can be examined component by component: Are the revenue, earnings, and cash flow forecasts realistic? Is the assumed multiple consistent with the growth, returns on capital, and competitive position expected at that point? Are the key assumptions internally coherent? These questions are concrete and falsifiable. By contrast, the key assumptions in a DCF—particularly those embedded in the TV—often concern cash flows many years beyond any reasonable forecasting horizon and are rarely examined with the same rigor.

IV. A Hybrid Discipline

The strongest approach, in my view, is to use each tool for the purpose it serves best rather than forcing any single method to carry the full burden of setting a price target. Forward multiples are the primary instrument for establishing a price target because they align with how markets actually discover prices. DCF analysis is more useful as a diagnostic tool—for mapping reinvestment dynamics, testing the internal consistency of long-term assumptions, and building intuition about what a business must achieve to justify its current valuation. A third lens, private market value (PMV), can provide an independent cross-check.

In practice, building a price target using forward multiples begins with a detailed, driver-based forecast of revenues, margins, capital intensity, and profits through the target date. The analyst then applies a multiple to projected earnings—informed by the expected growth rate, returns on capital, competitive position, and risk profile at that horizon—and divides by diluted shares outstanding to arrive at a per-share price target.

The multiple itself can be further grounded by examining how comparable companies and the company itself have traded when their characteristics were similar.

Running a DCF in parallel serves a narrower purpose. It is used not to generate the price target, but to test whether the assumptions embedded in the multiple-based target are consistent with the business's long-term economics. When the DCF implies materially different long-term growth or returns on capital than those reflected in the exit multiple, the divergence itself becomes useful information. It may indicate that the exit multiple is too optimistic or too conservative, or that important elements of the business model have not been fully captured.

A third check—estimating the company's private market value (PMV)—can serve as a useful cross-check for any business, but it is especially powerful for companies with distinct operating segments. In essence, PMV reflects what a well-informed buyer would pay to acquire control of the entire business or any of its subsidiaries. This is typically estimated by valuing each business line separately by applying an appropriate multiple to its cash flows—multiples usually derived from comparable mergers and acquisitions (M&A) transactions involving similar businesses—and then adjusting for debt and cash on the balance sheet to arrive at an equity value. Like forward multiples, PMV analysis can also be projected forward to a future date by forecasting the cash flows of each business line and applying an appropriate exit multiple at that horizon.

When the gap between a company's PMV and its public market price becomes sufficiently wide, the spread itself can act as a catalyst. A material discount to private market value often attracts acquirers or activist investors, creating economic pressure that can pull the public stock price higher over time through anticipated corporate actions such as asset sales, spin-offs, or a change of control. In this context, PMV serves a dual role within the hybrid framework: it functions as an independent cross-check alongside forward multiples and DCF analysis, and it can also signal situations in which fundamental value has the potential to catalyze real upward pressure on the current market price.

The value of this approach lies not in achieving false precision, but in surfacing the assumptions that matter most. When the three tools converge, they reinforce confidence in the price target. When they diverge—particularly when PMV reveals a wide gap to the current market price—the tension often reveals both analytical inconsistencies and the potential for fundamental forces to narrow that gap over time.

V. Probabilistic Thinking Across Scenarios

Neither framework alone produces a point estimate that deserves a high degree of confidence, and this is worth stating plainly. While forward multiples offer a more practical and transparent approach than DCF models for setting price targets, even a well-constructed multiple-based target does not justify a high level of confidence in any single outcome. The appropriate response is therefore not to chase a level of precision that investing—unlike physics—cannot realistically deliver, but rather to build probabilistic structure around the analysis. This means developing a distribution of outcomes rather than anchoring on a single central case, assigning explicit probabilities to scenarios rather than relying on vague qualitative labels, and sizing positions according to both expected return and the shape of the distribution.

A practical structure begins with three scenarios—base, bull, and bear—each built on explicit assumptions about competitive trajectory, revenue growth, margin evolution, reinvestment intensity, and exit multiple. The base case represents the analyst’s best estimate of the most likely path. The bull case captures a credible upside in which competitive advantages prove more durable or execution exceeds expectations. The bear case captures a credible downside in which competitive pressure intensifies, reinvestment returns disappoint, or multiple compression amplifies the fundamental shortfall. These scenarios should be grounded in specific business outcomes rather than broad macro assumptions.

The architecture of these scenarios differs meaningfully by business type. For businesses whose value resides primarily in the balance sheet or in current earnings power, the bear case is typically anchored in observable competitive pressure, margin compression, or cyclical demand deterioration. The feedback is relatively prompt—not unlike ordering a steak at your favorite restaurant that you know well. You have a prior, the result arrives quickly, and you update your view on the spot.

For businesses whose value resides in network-reinforced platforms or foundational intelligence capabilities—the frontier of 4.0—the meal takes years to arrive and you are not entirely sure what you ordered. The bear case is no longer primarily about near-term earnings disappointment. It is about whether the moat being built is real: whether the data flywheel accelerates, whether the network effect is defensible, and whether the inference cost curve improves fast enough to matter. These are questions about the durability of an advantage that is not yet fully visible in reported earnings.

Because accounting in these businesses often understates the intrinsic value being built, the bear case requires the analyst to distinguish between two very different things: earnings pressure that reflects deliberate reinvestment in a compounding asset, and earnings pressure that reflects a business model that is not working. Conflating the two remains one of the most common and costly errors when valuing this generation of compounders.

Scenario weights are explicit probability assignments. The probability-weighted price target represents the expected value of outcomes across scenarios. But the target alone understates the relevant information. An investment with a scenario-weighted target implying a 15 percent five-year compound return but a bear case implying permanent capital loss deserves a fundamentally different position size—or none at all—than one with the same expected return and a bear case of modest underperformance. Expected value and variance are both inputs to a rational decision. Asymmetry—the relationship between the size of gains in the bull case and losses in the bear case—matters as much as the central estimate.

For Value Investing 4.0 businesses in particular, where accounting distortions can obscure early signs of fundamental progress or deterioration, getting asymmetry right is not a refinement. It is the mechanism by which a margin of safety is maintained dynamically rather than established once at entry. The appropriate response to compounding uncertainty is therefore not avoidance but calibration: applying a stricter margin of safety, taking a smaller initial position, and identifying in advance the specific leading indicators—Engagement Value per Unit Time, cohort-level unit economics, reinvestment returns, talent retention, and inference cost trajectory—that would cause a rational analyst to revise scenario probabilities.

Base-rate success probabilities decline as the time horizon of value creation extends—not because the businesses are worse, but because high-quality confirmation arrives later and with more noise. As a result, updating one's views as new evidence emerges becomes just as important as forming the initial estimate. A position entered on a specific competitive thesis should be actively interrogated as new data arrives: Are key metrics evolving consistently with the base case? Are the leading indicators tracking the projected path? Unanticipated developments should trigger explicit reassessment of scenario probabilities rather than rationalization. The real discipline of probabilistic thinking lies not in constructing scenarios, but in specifying in advance what evidence would change them—and then having the intellectual honesty to follow that evidence when it arrives.

VI. Implications for Value Investing 4.0

The framework I have described above is not a departure from value investing—it is value investing adapted to the era in which we operate. The core commitments remain unchanged: understand business economics at a fundamental level, insist on a margin of safety, think probabilistically about outcomes, and size positions according to the quality of the opportunity and the conviction of the analysis. What changes

across eras is the dominant mode and timing of value creation and, correspondingly, the tools best suited to measuring it.

In the current era, competitive advantages increasingly reside in network-reinforced platforms, proprietary data, and foundational intelligence capabilities. Two accounting distortions are particularly relevant to businesses building these advantages. First, reported earnings systematically understate intrinsic value, as the income statement expenses investments in platforms, data infrastructure, and talent that economic logic would capitalize. Second, a DCF calibrated to today's suppressed profits often fails to fully capture the normalized earning power that aggressive reinvestment is building. While an analyst could theoretically adjust future cash flows and the TV to reflect declining reinvestment intensity, in practice, these adjustments regarding timing and the extent of cash flow normalization are difficult to make with confidence. As a result, DCF models frequently undervalue the companies most deliberately compounding long-run value.

The approach of projecting the business's state in the future through driver-level analysis and applying a multiple appropriate to that future state helps address both distortions. It recognizes latent value being built today without mistaking temporary earnings pressure for permanent impairment. Running a DCF alongside the multiple-based target can serve as a useful check on the long-term consistency of the assumptions, but it is *not required* to generate the price target itself. This framework anchors the target in observable market pricing mechanisms while preserving fundamental discipline.

One further dimension deserves attention: the nature of the moat itself is changing. Physical distribution networks, brand loyalty, and regulatory licensing—the classic durable advantages that I still respect enormously—remain relevant but are increasingly joined and sometimes displaced by advantages rooted in data accumulation, model quality, inference cost curves, and feedback loop velocity. These advantages compound differently, erode differently, and require a different vocabulary of leading indicators to evaluate.

As David Reichberg has written in "Form Factor Risk to the Competitive Positioning of a Business," even established competitive positions can be disrupted when new technologies fundamentally change the way value is delivered or consumed.[7]

Reichberg's central insight is that the analyst's task is not to forecast the next device, but to distinguish advantages that are genuinely durable—retained across any medium—from those that are merely borrowed from the current one. A business whose competitive position depends on the conditioning loops of today's smartphone interface is exposed in ways that neither its user count nor its brand recognition will reveal in advance. The investor still reasoning from yesterday's competitive framework will systematically misprice the era's most durable compounders. Seeing those businesses clearly requires not just better models but better questions.

Conclusion: The Water, Examined

Wallace's image lingers because the most dangerous assumptions are not the ones we consciously adopt, but the ones we eventually stop seeing. For investors long trained on DCF models, the perpetuity assumption has functioned exactly that way: structurally essential to the model, rarely examined, and increasingly detached from a world in which competitive advantages erode more quickly than they once did.

---

[7] David Reichberg, "Form Factor Risk to the Competitive Positioning of a Business," *Tsai Capital*, May 2026, https://tsaicapital.com/form-factor-risk.php

To examine that water is not to abandon DCF analysis. It is to use the tool for what it does well—building disciplined intuition about business economics, stress-testing long-run assumptions, and maintaining clarity in the relationship between price and value. It is also to stop using it for what it cannot reliably do: generate price targets whose apparent precision rests on assumptions about cash flows many years beyond any reasonable forecasting horizon. The compound return an investor earns over any finite horizon is determined by the gap between the entry and target prices, not by the discount rate embedded in a model designed to value a business in perpetuity. These are different questions, and conflating them produces systematic confusion in both target-setting and portfolio allocation.

Forward earnings multiples, grounded in the same fundamental drivers that any serious valuation requires but applied to the expected state of the business at a defined future point, offer a more honest and practical instrument for setting price targets. They align with how markets actually discover prices. They are decomposable, falsifiable, and explicit about the judgments they require. When combined with disciplined scenario analysis, probabilistic thinking, and a margin of safety appropriate to the uncertainty of the horizon, they form a framework that is both intellectually rigorous and genuinely useful in practice.

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 July 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.