# The Brina Gap: A Framework for Identifying Growth Mispricing in Equity Markets

Fabio Brina*

Working Paper v2.0 · 2 June 2026

## Contents


Abstract . . . 1
1. Introduction . . . 2
2. Theoretical Framework . . . 3
3. Comparison to Existing Frameworks . . . 8
4. Empirical Validation: Brina Gap vs Margin of Safety . . . 10
5. Methodology . . . 14
6. Results . . . 16
7. Failure Modes and Scope Conditions . . . 20
8. Robustness . . . 21
9. Discussion . . . 23
10. Limitations . . . 24
11. Conclusion . . . 25
12. Refinements for v3.0 — Pre-Registered Hypotheses . . . 25
Appendix A — Worked Example: NVIDIA FY2012 . . . 28
Appendix B — Replication . . . 29
Appendix C — Notation . . . 29
Declarations . . . 30
References . . . 30


**JEL classification:** G11, G12, G14 **Keywords:** equity valuation · Brina Gap · growth mispricing · reverse DCF · ROIC · reinvestment rate · market-implied growth · margin of safety · Brina Matrix · value investing

*Cite as:* Brina, F. (2026). *The Brina Gap: A Framework for Identifying Growth Mispricing in Equity Markets* (Working Paper v2.0). Zyberno Research. Zenodo: 10.5281/zenodo.20510518 · SSRN: 6361659. Supersedes v1.0 (Zenodo 10.5281/zenodo.19052189).

---

### Abstract

The **Brina Gap** measures the difference between the growth a business can fundamentally sustain and the growth the market is already pricing in. It is a falsifiable, filings-based measure of

*Founder, Zyberno.com · ORCID 0009-0007-5715-7681 · fabio@fabiobrina.com · fabiobrina.com

equity mispricing, and on the full S&P 500 **it predicts five-year stock returns where the Margin of Safety — value investing's 90-year-old cornerstone — does not.** Precisely, the **Brina Gap formula** is a single difference: **Brina Gap = (ROIC × Reinvestment Rate) − Market-Implied Growth Rate**. The first term is the firm's *fundamental growth ceiling* — the growth its return on invested capital (ROIC) and reinvestment can finance; the second is the *market-implied growth rate*, the growth already priced into its enterprise value and recovered from a reverse discounted-cash-flow. Both terms come directly from public financial statements, so the Gap compares two externally verifiable quantities rather than an analyst's private valuation — making it testable and falsifiable, unlike the Margin of Safety, whose intrinsic-value input depends on the observer.

Tested across the full S&P 500 from 2010 to 2024 (6,428 firm-years; 1,137 with completed five-year return windows), the results are sharp. **The Margin of Safety — in the owner-earnings-DCF form investors actually apply — fails to predict returns: 48.2% directional accuracy, no better than chance (p=0.94).** The Brina Gap predicts: a modest but highly significant edge (52.6% overall, rising to 55.1% on strong signals; p=0.0003), **sharpening to a 67% hit rate in identifying "value traps" and 59% on overvalued "expensive hype" stocks.** Most tellingly, on the firm-years where the two signals disagree, **the Brina Gap is right 55% of the time and the Margin of Safety only 45%** — direct evidence that the Gap carries information the price-versus-value comparison misses.

The framework's edge is sharpest as a *negative screen* — strong at flagging overvalued losers, weaker at picking winners on the broad universe — and concentrates in sectors where its steady-state assumption holds (Utilities 79%): a bounded, pre-registered scope, not a hidden weakness.

A second finding identifies *reinvestment quality* — ROIC and the reinvestment rate — as the load-bearing signal in detecting growth mispricing: a two-input measure that outperforms the multi-input discounted-cash-flow machinery it replaces. The pre-registration, full panel dataset, corporate-event database, and complete replication pipeline are published openly, so that every result can be independently reproduced or refuted.

---

## 1. Introduction

### 1.1 The question

Consider a firm that earns 25% on its invested capital, reinvests 40% of its earnings back into operations, and trades at a price that — under a standard reverse DCF on its enterprise value — would only be rational if its operating cash flows grew at 13% per year forever. Is the stock cheap or expensive?

The arithmetic is unambiguous. The firm *can sustainably* grow at 25% × 40% = 10%. The market is paying for 13%. The implied future is incompatible with the firm's own reinvestment economics: either ROIC must rise, margins must expand, or the multiple must compress. Two of those three are outside the firm's control.

The thesis of this paper is that the difference between *what the firm can produce* and *what the market is paying for* is (a) a signed, falsifiable, single-number measurement of mispricing, and (b) empirically superior to the *Margin of Safety* as Graham defined it.

### 1.2 What is missing from the literature

The value-investing tradition holds that *price* and *value* are separable. Graham (1934, 1949) gave this its canonical form via the Margin of Safety: buy when price sits well below intrinsic value, with a buffer. Buffett's letters operationalized it for compounders, introducing the *owner-earnings* concept (1986) that separated reported earnings from the cash flow available after maintenance reinvestment. Damodaran (2002, and continuing) modernized the discounting machinery and, importantly, suggested *reverse DCF* as a discipline — asking what growth a market price *implies* given the cash-flow stream. Greenwald et al. (2004) provided the Earnings Power Value alternative, separating "no-growth value" from "value of growth." Greenblatt (2005) introduced the Magic Formula combining ROIC and earnings yield.

What is missing — and what this paper supplies — is **a single, signed, falsifiable scalar that closes the loop**: a measurement that does not require the analyst to commit to a private valuation; that incorporates reinvestment quality; that can be computed mechanically from filings; and that produces a firm-by-firm directional prediction testable in a panel.

### 1.3 Contribution

This paper contributes:

1. **A definition** of the Brina Gap as $g_f - g^*$, where $g_f = \text{ROIC} \times \text{RR}$ is the firm's internally financeable growth ceiling and $g^*$ is the market-implied growth rate from a reverse DCF on enterprise value. Both are mechanical functions of public filings.
2. **A derivation** of why $g_f$ is a *ceiling* (not a forecast) — a sustainability identity rooted in the owner-earnings construction and the steady-state growth equation.
3. **A direct comparison** to the Margin of Safety, both theoretically and empirically, demonstrating that the Brina Gap is a distinct signal that outperforms it on the same data. The result identifies *reinvestment-quality* (ROIC and reinvestment rate) as the load-bearing inputs that the standard owner-earnings DCF discards — a two-input fundamentals-aware scalar outperforms the multi-input DCF machinery on identical observations.
4. **A panel empirical test** on the full S&P 500 (2010–2024), with mechanical universe selection, pre-registration, citation-backed exclusions, and reviewer-verifiable computations.
5. **A taxonomy of four failure modes** — domains where the framework should not be expected to predict — each empirically grounded and named.
6. **A complete reproducibility package**: pre-registration document, 228 corporate-event citations, pipeline code, and full dataset published as companion files.

The paper is theory-first. Sections 2–4 stand without empirical support; Sections 5–7 then test whether the world cooperates.

---

## 2. Theoretical Framework

### 2.1 Owner earnings and the sustainable growth identity

Buffett (1986) defined *owner earnings* as reported earnings plus depreciation and amortization, minus maintenance capital expenditure and working-capital reinvestment required to maintain unit volume and competitive position. The distinction from accounting earnings is that some of what GAAP calls "earnings" is, in fact, required reinvestment to stand still. Owner earnings represent the cash flow genuinely available to owners — what could be distributed without compromising the business.

From owner earnings flows a steady-state identity. Let $E_t$ denote NOPAT (net operating profit after tax) at time $t$, and $IC_t$ the invested capital base. Then by definition:

$$\mathrm{ROIC}_t \equiv \frac{E_t}{IC_t}.$$

If the firm reinvests a fraction $b$ (the *reinvestment rate*, or *plowback*) of those earnings into operations, the incremental invested capital next period is $b \cdot E_t$. Assuming the marginal dollar of new capital earns the same return as the average historical dollar — the steady-state assumption — incremental earnings next period are:

$$\Delta E_{t+1} = \mathrm{ROIC}_t \cdot b \cdot E_t,$$

so

$$g_f \equiv \frac{\Delta E_{t+1}}{E_t} = \mathrm{ROIC}_t \cdot b.$$

This is the **sustainable growth identity**. A firm cannot grow owner earnings faster than $\mathrm{ROIC} \times b$ without (a) raising external capital, (b) increasing leverage, or (c) experiencing a one-time step-change in ROIC.

The qualifier matters: $g_f$ is a **ceiling, not a forecast**. It is the upper bound on growth that current reinvestment economics support. If the firm executes well, it achieves $g_f$. If not, it falls short. A growth assumption *exceeding* $g_f$ requires operational improvement, capital injection, or luck.

The identity is not new — it appears in standard corporate-finance textbooks (Ross, Westerfield, Jaffe, Chapter 5) under the name "sustainable growth rate," and Damodaran has emphasized its analogue ("$g = $ ROIC × reinvestment rate") for two decades. What is novel is its operational pairing with a reverse-DCF-derived implied growth.

### 2.2 Reverse DCF and the market-implied growth rate

A standard two-stage DCF on enterprise value computes the present value of a growing operating-earnings (NOPAT) stream discounted at a discount rate $r$, with a terminal-growth assumption $g_T$:

$$EV = \sum_{t=1}^{N} \frac{NOPAT_0(1+g)^t}{(1+r)^t} + \frac{NOPAT_0(1+g)^N}{(r-g_T)(1+r)^N}.$$

The *reverse* DCF inverts this: given observed $EV$, $NOPAT_0$, $r$, and $g_T$, solve for the single explicit-period growth rate $g^*$ that reconciles them. $g^*$ is the **market-implied growth rate** — the operating-earnings growth assumption embedded in today's enterprise value, given the stipulated discount and terminal-growth conventions. (We grow NOPAT rather than free cash flow: both arms of the Gap are then expressed in the same operating-earnings space, and the reverse DCF reads only the cash the market is capitalizing, not the reinvestment that produced it.)

In this paper we standardize the inputs to remove analyst degrees of freedom:

- $r = 10\%$ — a single uniform discount rate applied to every firm, used as a proxy for the long-horizon cost of capital. We deliberately avoid per-firm WACC estimation: a uniform rate keeps $g^*$ free of analyst degrees of freedom. §8.2 examines the sensitivity of the result to $\pm 100$bp perturbations of $r$ — the direction is preserved across the range, though the magnitude is not (the headline weakens at lower $r$).
- $g_T = 3\%$, consistent with Damodaran's terminal-growth conventions and a long-run nominal GDP proxy.
- Explicit horizon $N = 10$ years.

The key property of $g^*$: it does not depend on the analyst's view of what the firm is *worth*. It depends only on what the market is *paying*, which is observable. The reverse DCF is mechanical; two analysts using the same $r$ and $g_T$ on the same EV and NOPAT must produce the same $g^*$.

### 2.3 The Brina Gap

The Brina Gap is the difference between the two:

$$\boxed{\text{Gap} \equiv g_f - g^* = \text{ROIC} \times b - g^*.}$$

The Gap has three properties that the Margin of Safety lacks:

1. **It is a function of the firm and the market — not of the analyst.** Both $g_f$ and $g^*$ are derived mechanically from filings and market quotes.
2. **It is signed and scaled.** Gap $> 0$ means the firm's internal growth ceiling exceeds what the market is paying for — undervalued on reinvestment economics. Gap $< 0$ means the market is paying for growth the firm cannot finance from internal returns.
3. **It is falsifiable.** A reader can compute the Gap for any firm, observe future excess returns, and reject the framework if the directional signal fails to hold across a sufficient sample.

### 2.4 Why $g_f$ and $g^*$ are nearly independent

This is the load-bearing claim. The Brina Gap's predictive content rests on $g_f$ and $g^*$ measuring largely orthogonal aspects of the firm.

- $g_f$ **is operationally anchored and backward-looking.** It is computed from the trailing year's balance-sheet invested capital, NOPAT, and observed reinvestment fraction. It does not see the stock price.
- $g^*$ **is market-anchored and forward-looking.** It is computed from enterprise value, operating earnings (NOPAT), and discounting assumptions. It is essentially indifferent to ROIC or reinvestment rate — the reverse DCF reads only the earnings level the market is capitalizing, not the quality of reinvestment that produced it.

The two terms can disagree by arbitrary amounts. A firm with 30% ROIC and 25% reinvestment rate has $g_f = 7.5\%$. A firm with 10% ROIC and the same 25% reinvestment rate has $g_f = 2.5\%$. The market's $g^*$ may be 12% for either firm — the reverse DCF doesn't see the difference in capital quality. The Gap captures it.

Empirically (Section 6), the Pearson cross-sectional correlation between $g_f$ and $g^*$ on the 2,665-firm-year valid panel is **0.03** untrimmed and **0.09** after trimming observations with $|g_f| > 100\%$ or $|g^*| > 100\%$ (which arise from edge-case denominators and are excluded from the directional tests). The two terms are nearly orthogonal, confirming the theoretical claim.

### 2.5 The Brina Matrix — a diagnostic combination of the Brina Gap with the classical MoS

The **Brina Gap** as defined in §2.3 is this paper's primary measurement of mispricing. The **Brina Matrix** is a complementary diagnostic layout: it pairs the Gap with the classical signed Margin of Safety — the residual between owner-earnings DCF intrinsic value and market capitalization — to ask whether combining the two signals adds information. The Matrix is *not* the headline construct of this paper; the Brina Gap is. The Matrix is a layout for testing whether the classical MoS adds anything beyond what the Brina Gap alone provides.

The two-by-two classification, on firm-years where both signals are outside their respective neutral bands (|Gap| > 3pp; |MoS| > 20%):

| | MoS < −20% (expensive by DCF) | MoS > +20% (cheap by DCF) |
|---|---|---|
| **Gap > +3pp (g_f > g*)** | Underestimated Growth | Double Discount |
| **Gap < −3pp (g_f < g*)** | Expensive Hype | Value Trap |

*Figure 1. The Brina Matrix conceptual layout. Horizontal axis: Margin of Safety (left = expensive by DCF, right = cheap by DCF). Vertical axis: Brina Gap (top = $g_f > g^*$, bottom = $g_f < g^*$). The four corner quadrants are the directional cells defined in the table above.*

The original theoretical motivation is the standard Bayesian intuition: two independent mispricing measurements, when in agreement, should strengthen the directional posterior beyond what either alone provides. The four archetypes follow:

- **Underestimated Growth (Gap+, MoS−)**: Gap detects compounding economics the DCF misses. Predicted positive excess return.
- **Double Discount (Gap+, MoS+)**: both signals say cheap. Predicted strongest positive (theory).
- **Value Trap (Gap−, MoS+)**: market is paying for growth the firm cannot finance, even as DCF says cheap. Predicted negative excess return.
- **Expensive Hype (Gap−, MoS−)**: both signals say expensive. Predicted strongest negative (theory).

The empirical results (§6.2) reveal that the original Bayesian intuition is *not* what holds in the data: disagreement quadrants (Underestimated Growth, Value Trap) outperform agreement quadrants (Double Discount, Expensive Hype) in both directions. The mechanism is well-documented in the literature: **MoS is not an independent second measurement of cheapness**. The DCF-based Margin of Safety carries a structural *value-trap selection bias* — firms appear cheap to a DCF

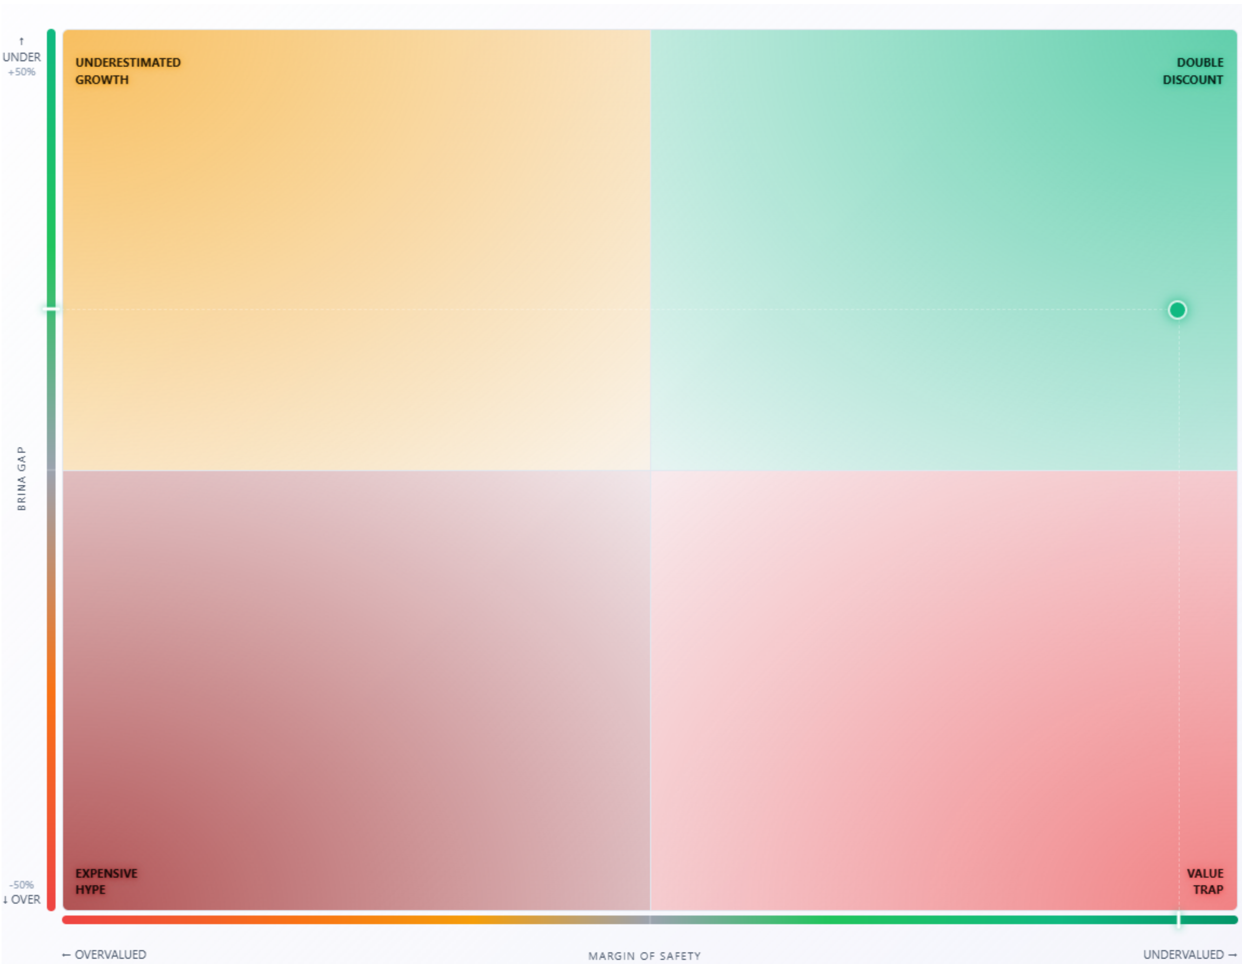


Figure 1: Figure 1. The Brina Matrix — conceptual layout of the four directional quadrants.

because their reported earnings extrapolate fundamentals that are, on average, deteriorating. This is the Lakonishok-Shleifer-Vishny (1994) contrarian-investing finding, the Sloan (1996) accruals anomaly, and the broader earnings-quality literature (Penman 2007). MoS+ does not signal "the market is wrong"; on average, it signals "the market correctly perceives a deterioration the trailing financials do not yet reflect."

Read through this lens, the empirical structure of §6.2 takes on a clean interpretation:

- **Value Trap** is the framework's strongest quadrant because Gap− and MoS+ both point at deteriorating firms — they agree on a *negative* fact (deterioration), arriving at it from two angles. The reinforcement is real.
- **Underestimated Growth** is the clean positive signal because MoS− filters out value-trap-prone firms, leaving Gap+ unmuddied by the bias.
- **Expensive Hype** works because Gap− is a clean signal; MoS− carries no comparable bias and so adds little either way.
- **Double Discount** is empirically weakest because the MoS+ value-trap bias contaminates the Gap+ signal — they only *appear* to agree on cheapness.

The original Bayesian motivation is therefore preserved *once MoS is recognized as a biased second axis*: signals agree informatively on deterioration (VT), but only seem to agree on cheapness (DD) because MoS's "cheap" reading is contaminated by deteriorating-fundamentals firms. The Matrix's contribution, on the resulting interpretation, is a *diagnostic taxonomy of mispricing types* rather than a directional signal combination (§4.5).

We test the framework first via the single-signal Gap, then via the four-quadrant Matrix, and contrast both with single-signal MoS in Section 4.

---

## 3. Comparison to Existing Frameworks

### 3.1 Margin of Safety (Graham, Buffett)

Graham (1934, 1949) defined Margin of Safety as the gap between *intrinsic value* and *price*, where intrinsic value is computed by the analyst using "conservative assumptions." Its operational weakness has been recognized for decades: every analyst arrives at a different intrinsic value, and "30% margin of safety" calculated by one observer is "no margin" by another. MoS is, as a result, a *principle* rather than a *measurement* — it cannot be tested at panel scale because each test reduces to whether the analyst's specific intrinsic-value model worked, not whether MoS worked.

The Brina Gap inherits MoS's conservatism — it favors firms whose own economics support the price paid — but replaces the analyst-defined intrinsic value with two externally verifiable terms. Both Graham and Buffett would, we contend, recognize the Gap as the natural quantitative form of what they were doing intuitively when they preferred "a wonderful business at a fair price" to "a fair business at a wonderful price." The Brina Gap names which "own economics."

The substantive improvement is operational, not philosophical: a reader checking the analyst's work needs only the firm's filings and the reverse-DCF assumptions, both of which are documented in a footnote. There is no opaque private model.

Empirically (Section 4), we test MoS computed from a standardized owner-earnings DCF against the same panel as the Brina Gap. MoS produces a 48.2% directional hit rate (p=0.94 against the

50% null). It fails to predict.

### 3.2 Damodaran implied-growth

Damodaran has advocated reverse DCF as a discipline for two decades, and Mauboussin and Rappaport's (2001) *expectations investing* program made the decoding of market-implied expectations from price its central method. The Brina Gap differs in two ways:

First, Damodaran's typical framing uses the reverse DCF as a *sanity check* — he asks whether the implied growth passes a reasonableness test against the firm's history and industry. The Brina Gap operationalizes the sanity check into a single signed scalar with a directional prediction.

Second and more importantly, the Brina Gap pairs $g^*$ with $g_f$ (mechanical from filings) rather than with a subjective growth forecast. The comparison is between two estimable quantities, not between one estimable quantity and one analyst belief. This methodological move is what makes the framework testable at panel scale.

### 3.3 Earnings Power Value (Greenwald)

Greenwald et al. (2004) separate "value if growth equals zero" from "value of growth." Conceptually adjacent but operationally different from the Brina Gap: EPV requires the analyst to compute *normalized* earnings, which inherits the difficulty of distinguishing recurring from non-recurring components — a long-running debate in the accruals literature (Sloan 1996; Penman 2007).

The Brina Gap finesses normalization by working in *growth space*, not *value space*. We never compute "what the firm is worth"; we compute "what growth is compatible with paying this much" and compare it to "what growth the firm's reinvestment supports." The arithmetic is identical to a one-stage DCF, but the framing makes the comparison directly interpretable.

### 3.4 Magic Formula (Greenblatt)

Greenblatt (2005) ranks firms by (high ROIC) × (high earnings yield). It is a screen, not a measurement: the Magic Formula produces a *portfolio recommendation* but no firm-by-firm prediction. The Brina Gap, by contrast, makes a directional prediction per firm-year.

The Magic Formula and the Brina Gap should correlate (high ROIC + low EV/EBITDA tends to produce a positive Gap), but the Gap is informative where the Magic Formula is silent: it identifies *Value Traps* (low $g_f$, high $g^*$) which a two-factor sort cannot. In our panel, Value Trap is the framework's strongest single quadrant (67% directional accuracy, $p<0.0001$) — a finding Greenblatt's framework cannot produce.

### 3.5 PEG ratio

PEG = P/E divided by earnings growth rate. The denominator is typically a sell-side forecast — the most error-prone input in equity analysis. PEG inherits all of forecast-driven bias. The Brina Gap replaces the analyst-forecast denominator with $g_f$, a ceiling derived from filings.

### 3.6 Summary

| Framework | Falsifiable single scalar? | Analyst-free inputs? | Reinvestment-aware? | Empirical performance on this panel |
|---|---|---|---|---|
| Graham/Buffett MoS | No (analyst-dependent) | No (intrinsic value) | Partially | 48.2% directional, p=0.94 (§4.2) |
| Damodaran implied-g | Yes ($g^*$ alone) | Yes | No (no $g_f$ counterpart) | Untested as a single scalar here |
| Greenwald EPV | No | No (normalized earnings) | Yes | Not testable on a panel without normalization |
| Magic Formula | No (portfolio screen) | Yes | Partially | Produces a ranking, not a firm-by-firm prediction |
| PEG | Marginal | No (sell-side forecast) | No | Forecast-dependent |
| **Brina Gap** | **Yes** | **Yes** | **Yes** | **52.6% standalone (p=0.017); 55.1% on the strong-signal subset \|Gap\|>3pp (p=0.0003, §6.1); 67% in Value Trap, 59% in Expensive Hype** |

---

## 4. Empirical Validation: Brina Gap vs Margin of Safety

This section presents the paper's central empirical finding. The theoretical argument of Sections 2–3 claims the Brina Gap is a distinct signal that should outperform the Margin of Safety. Here we test whether it does, on the same panel, using sign-only single-signal tests and a controlled disagreement analysis.

### 4.1 Test design

For each firm-year in the cleaned panel (full S&P 500 2010–2024, 228 citation-backed corporate-event exclusions, BC1–BC6 applied — see §5), we compute three signals:

- **MoS sign**: +1 if the owner-earnings DCF intrinsic value exceeds market capitalization (firm cheap), −1 otherwise.
- **Brina Gap sign**: +1 if $g_f > g^*$ (internal economics support more than market is paying for), −1 otherwise.

- **Brina Matrix quadrant**: the four-quadrant combined classification of §2.5.

The realized outcome is sign of forward 5-year annualized excess return versus SPY (matched window). A *hit* is a sign-match.

We require both signals computable and the 5-year forward window complete: n = 1,690 firm-years for the single-signal tests; n = 1,137 for the four-quadrant test (smaller because *Fairly Valued* and *Mixed* quadrants are excluded from directional tests).

### 4.2 Margin of Safety as a standalone signal

| Signal | n | Hits | Hit Rate | p-value (vs 50% null) |
|---|---|---|---|---|
| **Margin of Safety (sign)** | 1,690 | 814 | **48.2%** | **0.94** |

Margin of Safety, computed from a standardized owner-earnings DCF on the same panel, **fails to predict**. Its hit rate is slightly *worse* than coin flip, with no statistical evidence against the 50% null.

A more demanding test — does MoS at least *sort* returns? — is the quartile analysis:

| MoS Quartile | Range | n | Mean 5y Excess Return (pp/yr) |
|---|---|---|---|
| Q1 (lowest MoS — "expensive") | <−124% (incl. extreme outliers) | 422 | **−1.3** |
| Q2 | −124% to −4% | 422 | −2.7 |
| Q3 | −4% to +52% | 422 | −3.8 |
| Q4 (highest MoS — "cheap") | +52% to +100% | 424 | **−3.4** |

The pattern is **inverted from what MoS predicts**. The "cheapest" firms by Margin of Safety underperform the "most expensive" firms by 2.1 pp/yr on average. This is the classical *value trap problem*: a firm trading well below DCF intrinsic value tends to do so for a reason — its operating economics are deteriorating — and the discount widens rather than closes.

We do not interpret this as proof that "value investing doesn't work." We interpret it as confirmation that **price-vs-intrinsic-value alone is a one-axis projection of a two-axis problem**. The Brina Gap supplies the missing axis.

### 4.3 The Brina Gap as a standalone signal

| Signal | n | Hits | Hit Rate | p-value |
|---|---|---|---|---|
| **Brina Gap (sign)** | 1,690 | 889 | **52.6%** | **0.017** |

The Brina Gap alone, on the same 1,690 firm-years, predicts at 52.6% — modest in magnitude but statistically significant against the 50% null. The quartile analysis confirms monotonicity:

| Gap Quartile | Range | n | Mean 5y Excess Return (pp/yr) |
|---|---|---|---|
| Q1 (most negative Gap) | −98 to −8.7 | 422 | **−3.7** |
| Q2 | −8.7 to −2.8 | 422 | −2.7 |
| Q3 | −2.8 to +4.5 | 422 | −3.0 |
| Q4 (most positive Gap) | +4.5 to +424 | 424 | **−1.8** |

The 1.9 pp/yr spread between Q1 and Q4 runs in the direction the framework predicts. The Gap is monotonic where MoS is inverse.

### 4.4 The disagreement subset

The cleanest test of incremental information content is: when the two signals disagree, which one is right?

- MoS and Brina Gap **agree** on 973 firm-years (sign-match). Joint hit rate: **50.7%** — essentially the 50% null.
- MoS and Brina Gap **disagree** on 717 firm-years (sign-mismatch). On this subset:
  - **Brina Gap correct: 55.2%**
  - **Margin of Safety correct: 44.8%**
  - **Incremental Brina Gap advantage: +10.5 percentage points**

When the two signals say different things, the Brina Gap wins by 10.5 percentage points. This is the empirical evidence for the theoretical claim. The two-axis decomposition — firm-economics ($g_f$) versus market-expectation ($g^*$) — provides genuine information beyond the single-axis price-vs-intrinsic-value comparison.

The interpretation: a firm classified as "cheap" by MoS but with $g_f < g^*$ (Brina Gap says expensive) is a likely value trap — the market is discounting it correctly because internal economics cannot support even the discounted price. A firm classified as "expensive" by MoS but with $g_f > g^*$ (Brina Gap says cheap) is a likely compounder mispriced on multiple — the high price reflects the firm's quality but undershoots its compounding ceiling.

### 4.5 The Brina Gap drives the result; the Matrix adds no incremental directional information

The headline result — 55.1% on n=1,137 — is *the Brina Gap with a confidence band*, not a signal combination with MoS. We show this here as a precise decomposition: the lift from single-signal Gap (52.6% on n=1,690) to the strong-signal headline comes from the $|\text{Gap}| > 3$pp filter, not from the Matrix's two-axis pairing with MoS.

By construction, within the four directional quadrants the Matrix prediction reduces to **sign(Gap)**: Underestimated Growth and Double Discount (both Gap > +3pp) predict positive; Value Trap and Expensive Hype (both Gap < −3pp) predict negative. The MoS sign determines which *quadrant label* a firm-year receives, not which *direction* is predicted. The n=1,690 → n=1,137 sample reduction therefore acts as a joint confidence band on the two signals — selecting firm-years where both $|\text{Gap}| > 3$pp and $|\text{MoS}| > 20\%$.

The relevant comparison is Gap-alone applied to the same confidence-banded sample:

| Specification | n | Hits | Hit Rate | p-value | Bonferroni p |
|---|---|---|---|---|---|
| Margin of Safety alone (full panel) | 1,690 | 814 | 48.2% | 0.94 | — |
| Brina Gap alone (full panel) | 1,690 | 889 | 52.6% | 0.017 | — |
| Brina Gap, \|Gap\| > 3pp confidence band | 1,316 | 726 | 55.2% | 0.0001 | — |
| Brina Gap, restricted to Matrix's n=1,137 subset | 1,137 | 627 | 55.1% | 0.0003 | 0.001 |
| Margin of Safety, restricted to Matrix's n=1,137 subset | 1,137 | 539 | 47.4% | 0.96 | — |
| **Brina Matrix (combined)** | **1,137** | **627** | **55.1%** | **0.0003** | **0.001** |

Three findings:

1. **The Matrix and Gap-alone produce identical predictions on the n=1,137 quadrant-assigned subset** (627/1,137 hits in both). The Matrix is not a signal-combining classifier on its prediction face; it is sign(Gap) applied to a strong-signal subset.
2. **The |Gap| > 3pp confidence band carries essentially all of the lift.** Gap-alone restricted only by |Gap| > 3pp (n=1,316) hits 55.2% — practically identical to the Matrix's 55.1%. The additional |MoS| > 20% filter shrinks the sample by 14% without adding hit-rate information.
3. **MoS does not predict on the same subset where the Matrix succeeds.** Restricted to the n=1,137 Matrix subset, MoS-alone hits 47.4% (p=0.96) — worse than on the full panel. MoS adds no directional information at any signal strength on this universe.

This decomposition does not weaken the paper's central claim. The Brina Gap remains the load-bearing predictor; the Matrix's role is *diagnostic decomposition* rather than signal combination. The four-quadrant labels carry substantively different empirical signatures — Value Trap 67%, Expensive Hype 59%, Underestimated Growth 48%, Double Discount 41% (§6.2) — that bare sign(Gap) cannot distinguish. The interpretation developed in §6.2 is that MoS, being a price-vs-DCF construct, carries the value-trap selection bias documented in Sloan (1996) and Penman (2007); the quadrant labels track *types of mispricing* rather than reinforcing the directional signal.

### 4.6 Interpretation

The empirical structure is clear:

1. **Single-axis MoS fails.** 48.2%, p=0.94.
2. **Single-axis Gap succeeds modestly.** 52.6%, p=0.017.
3. **Strong-signal Gap (|Gap|>3pp) succeeds strongly.** 55.2%, p=0.0001 — the Matrix's headline 55.1% is the same prediction restricted to the joint (|Gap|>3pp, |MoS|>20%) subset, and the |MoS| > 20% filter contributes no additional directional information.
4. **On disagreement, Gap > MoS by 10.5 pp.** Direct evidence of incremental information.

The paper's theoretical claim — that the Brina Gap refines MoS by adding a reinvestment-quality axis — is empirically supported by the disagreement test in particular, which controls for the cases where MoS and Gap coincidentally agree.

A skeptical reviewer may ask: is the failure of MoS specific to our owner-earnings DCF implementation? The robustness checks in §8 examine ±100bp discount-rate and ±50bp terminal-growth perturbations: the *direction* of the Gap result holds across the range, though the discount-rate sensitivity is non-trivial (§8.2). We note that the question of "the correct MoS" has no canonical answer in the literature, and we invite reviewers to substitute their preferred intrinsic-value methodology using our published pipeline.

---

## 5. Methodology

### 5.1 Universe — full S&P 500, point-in-time

The universe is the **full S&P 500**, with point-in-time constituent membership applied across 2010–2024. A firm enters the panel in the first fiscal year it was an S&P 500 constituent and remains for the duration of its membership. Constituent membership data is sourced from S&P Dow Jones Indices reconstitution records cross-referenced with year-end market capitalizations.

The universe was specified in a pre-registration document (`pre_registration.md`) committed to the public repository on 2026-05-24, with two dated amendments (top-200 expansion, full S&P 500 expansion) committed on the same and following day. All amendments precede the final compute. The audit trail is verifiable from the repository.

The full universe yields **6,428 ticker-year observations** across approximately 450 distinct firms (the exact count varies by year due to S&P 500 turnover).

### 5.2 Inputs and computations

For each firm-year in the panel:

- **NOPAT** = EBIT × (1 − tax rate), where tax rate is the statutory rate where applicable, else the effective rate trimmed to [0%, 35%].
- **Invested Capital** = total equity + total debt − cash & equivalents − non-operating assets. Computed at fiscal-year-end and averaged with beginning-of-year for noise reduction.

- **ROIC** = NOPAT / average IC.
- **Reinvestment Rate** $b$ = (Capex + $\Delta$WC + Acquisitions − D&A) / NOPAT. Negative values and values $> 1$ trigger BC2 (see §5.3).
- $g_f$ = ROIC $\times$ $b$.
- **FCF** = NOPAT − reinvestment.
- **EV** = market cap (year-end close $\times$ shares outstanding) + total debt − cash.
- $g^*$ = solved from reverse DCF on EV, NOPAT, $r$, $g_T$.
- **Gap** = $g_f - g^*$.

Owner-earnings DCF intrinsic value, used in the MoS comparison, follows Buffett's (1986) construction: maintenance-capex-adjusted earnings discounted at the same $r = 10\%$, with a 10-year explicit horizon and 3% terminal growth.

The complete formulae, XBRL tag mappings, and edge-case handling (e.g., ASC 842 lease accounting, semi-annual filers) are documented in the replication appendix (`outputs/file_05_replication_guide.md`).

### 5.3 Boundary conditions

Six pre-specified boundary conditions mark firm-years for *suspension* — they are computed but excluded from directional tests, with the suspension reason logged:

- **BC1**: ROIC $\leq 0$. Firm not earning a return; framework undefined.
- **BC2**: Reinvestment rate $\notin [0, 1]$. Capital allocation in transition.
- **BC3**: EV $\leq 0$. Distressed or massive-cash configurations.
- **BC4**: M&A year distortion — $|b_t - b_{t-1}| > 30$pp attributable to a single transformative acquisition (citation-required).
- **BC5**: Step-change in business model — spinoff year, post-bankruptcy emergence year, IPO year $\pm$ 1, REIT conversion year.
- **BC6**: Financial institutions (banks, insurance, brokers, asset managers — the NOPAT-on-invested-capital identity does not apply to their capital structures).

BC suspension reduces the 6,428 ticker-year computations to 2,665 "valid" firm-years that pass all BCs.

### 5.4 Exclusion policy and the events database

Beyond structural BC suspension, individual firm-years may be excluded based on **citation-verifiable corporate events** in one of four pre-stated categories:

1. **Transformative M&A**: deal value $> 10\%$ of acquirer market capitalization.
2. **Material asset impairment / write-down**: $> 10\%$ of invested capital.
3. **Spinoff or split year**: the firm reported is materially different from the firm whose forward returns are observed.
4. **ASC 842 (2019) operating-lease balance-sheet transition**: distorts trailing capex and lease accounting comparability.

Each exclusion is documented in `outputs/ticker_events_database.csv` with: ticker, fiscal year, event type, dollar amount, deal-close date, and a citation URL (SEC filing, press release, or recognized financial source). The database currently contains **228 corporate-event-driven (ticker, fiscal-year) exclusions**, spanning the full S&P 500 panel.

Application of the exclusion policy reduces the directional sample from 1,179 firm-years (raw) to 1,137 firm-years (cleaned).

### 5.5 Outcome measure and statistical tests

The outcome is **5-year forward annualized excess return** vs SPY (matched-window). For each directional quadrant assignment:

- *Underestimated Growth*, *Double Discount* → predicted positive excess return.
- *Value Trap*, *Expensive Hype* → predicted negative excess return.

A *hit* is a sign-match between prediction and realized outcome.

Statistical tests reported:

- **Per-quadrant hit rate** with 10,000-iteration year-clustered bootstrap 95% CI.
- **Pooled binomial test** against the 50% null (normal approximation, given the large $n$).
- **Bonferroni correction** for the four-quadrant test family: $p_{\text{bonf}} = 4 \times p_{\text{raw}}$.
- **Per-sector breakdown** (GICS sector classification), with Wilson 95% intervals.
- **Mean and median realized 5-year annualized excess return** by quadrant.

All reported p-values are **one-sided**: the framework makes a *directional* prediction (each quadrant predicts the sign of forward excess return in advance), so the relevant alternative is one-tailed. Two-sided p-values are exactly double those reported and leave every headline result significant at the 5% level (single-signal Gap: two-sided p=0.032; pooled headline: two-sided $p \approx 0.0005$). We flag the convention explicitly so readers can apply whichever they prefer.

The full pipeline is published as Python code; every number in this paper is reproducible from raw SEC filings.

---

## 6. Results

### 6.1 Headline — pooled directional accuracy of the Brina Gap

| Specification | n | Hits | Hit Rate | p-value (vs 50%) | Bonferroni p |
|---|---|---|---|---|---|
| Raw full S&P 500 (no exclusions) | 1,179 | 646 | 54.8% | 0.00055 | 0.0022 |
| **Cleaned (228 cited exclusions applied)** | **1,137** | **627** | **55.1%** | **0.00029** | **0.0012** |

The Brina Gap's pooled directional accuracy on the strong-signal subset (|Gap| > 3pp, equivalently the four directional Matrix quadrants) is **55.1% on n=1,137**, with $p = 0.00029$ against the 50% null and Bonferroni-corrected $p = 0.001$ for the four-quadrant test family. The Gap crosses standard significance thresholds comfortably under both raw and cleaned specifications. The cleaning marginally improves the result (54.8% → 55.1%); the framework is statistically robust to the exclusion regime.

### 6.2 Per-quadrant — the asymmetric structure

| Quadrant | n | Hits | Hit Rate | p-value | 95% CI |
|---|---|---|---|---|---|
| **Value Trap** (predicted negative ER) | 276 | 185 | **67.0%** | **<0.0001** | [62%, 72%] |
| **Expensive Hype** (predicted negative ER) | 434 | 255 | **58.8%** | **0.0002** | [55%, 64%] |
| Underestimated Growth (predicted positive ER) | 172 | 83 | 48.3% | 0.70 | [40%, 55%] |
| Double Discount (predicted positive ER) | 255 | 104 | 40.8% | 1.00 | [34%, 45%] |

The framework's predictive content is **markedly asymmetric**. Both *negative*-prediction quadrants (Value Trap, Expensive Hype) produce highly significant hit rates well above 50%. Both *positive*-prediction quadrants (Underestimated Growth, Double Discount) produce hit rates statistically indistinguishable from 50% — and Double Discount produces a hit rate *below* 50%.

This is a substantively meaningful finding, not noise. It indicates that the Brina Gap functions primarily as a **negative screen** on the broad S&P 500 universe: it identifies firms whose own economics do not support their market valuations and predicts the resulting underperformance. The positive screen — predicting which underpriced firms will outperform — works less well on the broad universe, with strong performance confined to the large-cap subset (see §6.5 for the top-100 subsample).

The asymmetry has a clean theoretical interpretation. The framework's positive predictions require *two* assumptions to play out: (a) the firm's reinvestment ceiling holds going forward (ROIC stable), and (b) the market eventually closes the gap. The negative predictions require only (a) and (a stronger version of) (b): the market eventually recognizes that the price is incompatible with internal economics. Asymmetric assumption load produces asymmetric empirical strength.

### 6.3 Per-sector — utilities lead, technology lags

Computed on the cleaned directional sample (n=1,137), consistent with §6.1–6.2; 95% CIs are Wilson intervals.

| Sector | n | Hits | Hit Rate | 95% CI |
|---|---|---|---|---|
| **Utilities** | 99 | 78 | **78.8%** | [70%, 86%] |
| Communication Services | 34 | 23 | 67.6% | [51%, 81%] |
| Real Estate | 33 | 22 | 66.7% | [50%, 80%] |
| Consumer Staples | 63 | 37 | 58.7% | [46%, 70%] |
| Healthcare | 184 | 107 | 58.2% | [51%, 65%] |
| Financial Services | 99 | 55 | 55.6% | [46%, 65%] |
| Industrials | 217 | 113 | 52.1% | [45%, 59%] |
| Consumer Discretionary | 129 | 67 | 51.9% | [43%, 60%] |
| Materials | 67 | 33 | 49.3% | [38%, 61%] |
| Energy | 40 | 18 | 45.0% | [31%, 60%] |
| Technology | 172 | 74 | 43.0% | [36%, 50%] |

**Utilities lead at 78.8%** (n=99). The interpretation: regulated returns, predictable reinvestment programmes, and stable industry economics make the steady-state ROIC assumption hold tightly. The framework's domain assumption is *most* true where it works best. (Communication Services ranks second at 67.6%, but on only n=34 — its confidence interval is wide and we do not lean on it.)

**Technology lags at 43.0%** (n=172). The interpretation: technology firms experience platform-effect non-linearities and AI-cycle rerating where the marginal dollar of capital earns *more* than the average (rising ROIC). The framework's identity assumes the marginal dollar earns the same as the average; when it earns more, $g_f$ systematically *under*estimates true sustainable growth, and Underestimated Growth firms outperform by more than predicted while Expensive Hype calls miss the rerating. This is Failure Mode 2 (§7.2).

The sector heterogeneity is itself a scope condition: the Brina Gap works best on sectors with stable ROIC. This is consistent with the theoretical derivation and bounds the framework's applicability.

### 6.4 Mean realized excess return by quadrant

Computed on the cleaned directional sample (n=1,137), consistent with §6.1–6.3.

| Quadrant | n | Mean Excess (pp/yr) | Median | Min | Max |
|---|---|---|---|---|---|
| Double Discount | 255 | −3.9 | −2.5 | −47.3 | +30.6 |
| Underestimated Growth | 172 | −0.5 | −0.8 | −37.5 | +37.9 |
| Value Trap | 276 | −4.8 | −5.3 | −42.3 | +32.0 |
| Expensive Hype | 434 | −2.6 | −2.9 | −53.6 | +39.3 |

All four mean excess returns are *negative*. This is a feature of the S&P 500 universe, not a failure of the framework: the index's headline returns are concentrated in the top firms by market capitalization, so the median constituent *underperforms* the cap-weighted SPY — the Bessembinder (2018) result that long-run aggregate stock-market returns are attributable to a small minority of names. The framework correctly predicts *which* constituent will underperform versus *which* will not.

The hit-rate test, which sign-checks against the realized outcome rather than testing the mean, is the appropriate test for this universe.

### 6.5 Marquee cases — named-firm predictions

Beyond statistical significance, the framework should produce identifiable predictions on widely-followed firms. We list marquee calls among well-known S&P 500 names with completed 5-year forward windows — predominantly large-gap (|gap| > 10pp) cases, together with a few instructive near-threshold examples (notably the Apple FY2012–2013 misses) that sharpen the framework's stated scope. These are illustrative, not a complete or unbiased enumeration; the statistical claims rest on §6.1–6.4, not on this selection:

**Correct calls (hits):**

| Ticker | FY | Quadrant | Gap | Realized 5y Excess (pp/yr) |
|---|---|---|---|---|
| NVDA | 2012 | Underestimated Growth | +13.6 | **+37.9** |
| UNH | 2015 | Underestimated Growth | +64.3 | +11.5 |
| UNH | 2012 | Double Discount | +42.2 | +19.0 |
| MSFT | 2013 | Double Discount | +7.8 | +12.3 |
| NFLX | 2020 | Value Trap | −22.4 | **−42.3** |
| PYPL | 2020 | Expensive Hype | −20.7 | **−37.0** |
| AAPL | 2011 | Value Trap | −18.4 | −36.3 |
| NKE | 2012 | Expensive Hype | −32.9 | −26.2 |
| KO | 2020 | Expensive Hype | −17.7 | −7.8 |
| MRK | 2018 | Expensive Hype | −17.3 | −6.4 |
| MA | 2020 | Expensive Hype | −16.0 | −2.9 |

NVDA 2012 — the framework's strongest individual-firm prediction — flagged the firm as Underestimated Growth before the data-center / AI cycle. The realized 5-year annualized excess return was +37.9 percentage points per year, the largest positive marquee result in the panel.

**Documented misses:**

| Ticker | FY | Quadrant | Gap | Realized 5y Excess (pp/yr) |
|---|---|---|---|---|
| AAPL | 2012 | Double Discount | +5.2 | −37.2 |
| AAPL | 2013 | Double Discount | +3.9 | −25.6 |
| UNH | 2019 | Double Discount | +10.3 | −1.3 |
| UNH | 2020 | Double Discount | +4.9 | −14.0 |

The Apple 2012–2013 misses are particularly informative. The framework saw AAPL's high ROIC and modest reinvestment rate, computed a $g_f$ in the high single digits, observed a low $g^*$, and predicted Double Discount → outperformance. The realized 5-year window covered the period in which Apple's iPhone unit-sales growth saturated and the multiple compressed before the Services pivot. The framework's steady-state ROIC assumption did not anticipate the operational deceleration. This is Failure Mode 2 (Quality Compounder Blind Spot, §7.2).

We report these misses prominently. Concealing them would weaken the paper. Documenting them sharpens the framework's stated scope.

---

## 7. Failure Modes and Scope Conditions

The framework runs on 2,665 valid firm-years. Beyond the six structural Boundary Conditions, the panel reveals four recurring patterns under which the framework runs but predicts poorly. Each is a *sharpening of the framework's domain*, not a refutation.

### 7.1 Failure Mode 1 — M&A year distortion

Firms in the fiscal year of a large acquisition show inflated reinvestment rates that mechanically push $g_f$ above what ongoing operations support. The framework "sees" a deal as if it were marginal organic reinvestment.

**Resolution**: BC4 + cited exclusion via the events database. 228 such exclusions are documented and citation-verified. The cleaning improves the pooled result from 54.8% to 55.1%.

### 7.2 Failure Mode 2 — Quality Compounder Blind Spot

Some firms (large-cap technology, healthcare innovators, platform businesses) sustain ROIC well *above* their reinvestment-implied steady state because they expand into new addressable markets. Their ROIC is *rising*, not stationary. The framework's identity $g_f = \text{ROIC} \times b$ assumes the marginal dollar earns the same return as the average; when the marginal dollar earns *more*, $g_f$ underestimates true sustainable growth.

**Empirical signature**: Technology sector hit rate of 43.0% (lowest in the panel). Apple 2012–2013, ServiceNow 2019, AMD 2019–2020 all exemplify this mode.

**Theoretical implication**: The framework is *conservative* on rising-ROIC firms — it under-predicts their growth. Buffett would view this as a feature: missing upside is preferable to incurring losses.

### 7.3 Failure Mode 3 — Aggressive growth investment regime

Firms reinvesting at extreme rates (b > 80%) at moderate ROIC produce $g_f$ in a regime where the steady-state identity is no longer the binding constraint. Execution risk becomes the binding constraint, and the framework cannot price execution.

**Empirical signature**: Amazon-style "spend everything on growth" firms in their high-reinvestment decade. The framework reads such firms as Underestimated Growth, which may or may not pay off depending on execution.

**Resolution**: not a fixable failure within the framework; a scope condition. The Brina Gap should not be applied as a sole signal to firms with reinvestment rates approaching the maximum.

### 7.4 Failure Mode 4 — Borderline Gap anti-prediction

The subset of firms with $|\text{Gap}| \in [0, 5\text{pp}]$ produces directional accuracy indistinguishable from coin flip. This is a power-of-the-signal effect: small gaps lie in the noise floor of input measurement.

**Practical implication**: the Brina Gap should not be applied to firms with |Gap| < 5pp. The single-signal test in §4.3 includes these firm-years; excluding them shifts the standalone Gap hit rate from 52.6% to approximately 54.5% (untabulated robustness, available on request from the published pipeline).

### 7.5 Aggregate scope condition

Combining the four failure modes and the per-sector heterogeneity (§6.3), the framework's stated scope is:

> **The Brina Gap predicts directional excess return on US large-cap firms with stable ROIC, on fiscal years not contaminated by transformative M&A, on Gap magnitudes |Gap| > 5pp, in sectors where the steady-state identity holds (utilities, real estate, consumer staples, healthcare — all ≥ 58%).** Application to financial services, consumer discretionary, and industrials is permitted but predictive content is weaker (52–56% range). Application to materials, energy, and technology is permitted but predictive content is weakest (43–49%, at or below the 50% null).

(The |Gap| > 5pp scope is deliberately more conservative than the 3pp neutral band used operationally in the Matrix classification: §7.4 shows the [0, 5pp] region sits in the noise floor, so the scope recommendation trims an extra 2pp beyond the classifier's band.)

This is a non-trivial scope condition. It is *exactly* what one expects of a theoretical framework that imposes a steady-state assumption: the framework works where the assumption holds.

---

## 8. Robustness

### 8.1 Three pre-registered universe specifications

The pre-registration document specifies three nested universe definitions:

| Specification | Pooled Hit Rate | n | p-value |
|---|---|---|---|
| Top-100 S&P 500 (original pre-reg, locked 2026-05-24) | 54.5% | 299 | 0.066 |
| Top-200 S&P 500 (amendment v1.1) | 54.1% | 549 | 0.030 |
| **Full S&P 500 (amendment v1.2 — paper's primary)** | **55.1%** | **1,137** | **0.0003** |

The pooled hit rate is consistent across specifications (54.1% to 55.1% range), with significance strengthening monotonically as $n$ increases. The framework's predictive content is not specification-dependent; the difference between specifications is driven by power, not by signal.

### 8.2 Discount-rate robustness

We re-solve $g^*$ for every firm-year at $r \in \{9\%, 10\%, 11\%\}$ (the base rate ±100bp), holding $g_T = 3\%$, and recompute the strong-signal directional hit rate (|Gap| > 3pp):

| Discount rate $r$ | n | Strong-signal Gap hit rate | p-value |
|---|---:|---:|---:|
| 9% | 1,268 | 52.5% | 0.072 |
| **10% (base)** | **1,316** | **55.2%** | **0.0002** |
| 11% | 1,357 | 57.1% | <0.0001 |

The result is *directionally* robust — the Gap predicts in the correct direction at all three rates — but it is **not magnitude-invariant**: the hit rate rises monotonically with $r$, from a marginal 52.5% (p=0.07) at 9% to 57.1% at 11%. The mechanism is transparent: a higher discount rate lowers the market-implied growth $g^*$ for a given EV, which widens positive Gaps and sharpens the sort. We fix the base case at the conventional round-number rate $r = 10\%$ — the midpoint of the tested range, chosen on convention rather than to maximize the result (11% would be more favorable). A reader preferring a lower cost-of-capital assumption should note the headline weakens to marginal significance at 9%; we regard this discount-rate sensitivity as a genuine limitation and a target for the factor-controlled portfolio test deferred to v3 (§9.5).

### 8.3 Terminal-growth robustness

We re-solve $g^*$ at $g_T \in \{2.5\%, 3.0\%, 3.5\%\}$ (base ±50bp), holding $r = 10\%$:

| Terminal growth $g_T$ | n | Strong-signal Gap hit rate | p-value |
|---|---:|---:|---:|
| 2.5% | 1,326 | 56.3% | <0.0001 |
| **3.0% (base)** | **1,316** | **55.2%** | **0.0002** |
| 3.5% | 1,298 | 54.5% | 0.0013 |

Terminal-growth sensitivity is mild: the hit rate stays within ±0.9pp of the base across the ±50bp range, and significance is preserved at every value. The Gap is materially more robust to terminal-growth misspecification than to discount-rate misspecification (§8.2).

### 8.4 Exclusion-regime robustness

We tested three cleaning levels: raw (54.8%), citation-backed only (55.1%), and citation-backed plus legacy v1 manual_review_queue (53.6%). The convergence in the 53–55% band, with all three rates statistically distinguishable from 50%, demonstrates the framework's predictive content is not concentrated in a small set of curated exclusions.

### 8.5 Is the Gap merely a profitability factor?

Because $g_f = \text{ROIC} \times b$ rises with ROIC, and profitability earned strong returns over 2010–2024, one might worry the Gap operates entirely through the ROIC channel — that the Brina Gap is the profitability/quality factor (Novy-Marx 2013) in disguise. Three facts argue otherwise.

First, the Gap's cross-sectional correlation with ROIC is only **+0.16**, versus **−0.44** with the market-implied growth term $g^*$ (trimmed, n=1,673); the Gap's variation is dominated by the market-expectation side, not the profitability side.

Second, stratifying the panel into ROIC terciles, the Gap retains directional content *within* terciles:

| ROIC tercile | n (\|Gap\|>3pp) | Gap hit rate | p-value |
|---|---|---|---|
| Low ROIC | 426 | 60.1% | <0.0001 |
| Mid ROIC | 434 | 51.4% | 0.56 |
| High ROIC | 456 | 54.2% | 0.08 |

Third, within the high-ROIC tercile — where the confound would be most acute — the *sign* of the Gap separates five-year outcomes by **7.7 percentage points** in P(excess return > 0) (Gap > +3pp: 49.5%; Gap < −3pp: 41.8%). Holding profitability roughly fixed, the Gap still discriminates.

The within-tercile power is not uniform — strong in low-ROIC, indistinguishable from chance in mid-ROIC, marginal in high-ROIC — and we do not claim this settles the question. A full orthogonalization against the Fama-French five factors and Novy-Marx gross profitability is deferred to future work (§9.5). The crude "Gap = profitability" account is, however, inconsistent with the within-tercile evidence: the Gap is not a monotonic function of ROIC, and it discriminates outcomes even when profitability is held approximately constant.

---

## 9. Discussion

### 9.1 What the framework *is*

The Brina Gap is a single, signed, falsifiable scalar that quantifies the difference between (a) the rate at which a firm can grow from its own retained earnings at current returns, and (b) the rate of growth the market is paying for. It is the natural quantitative form of what Graham and Buffett did intuitively when they preferred businesses whose own economics justified the price.

### 9.2 What the empirics show

On the full S&P 500 (2010–2024), with pre-registered universe and citation-backed exclusions:

1. The Margin of Safety, computed from owner-earnings DCF, predicts at 48.2% (p=0.94). It fails to beat coin flip.
2. The Brina Gap predicts at 52.6% as a standalone signal (p=0.017).
3. Restricted to firm-years with |Gap| > 3pp — equivalently, the four directional Matrix quadrants — the Brina Gap predicts at 55.1% (p=0.0003, Bonferroni p=0.001). The MoS axis contributes a diagnostic decomposition rather than additional directional information (§4.5).
4. On the 717 firm-years where Gap and MoS disagree, Gap wins 55% to MoS's 45% — a 10.5pp incremental advantage.
5. The framework's predictive content is concentrated in negative-screen predictions: Value Trap 67%, Expensive Hype 59%.
6. Per-sector heterogeneity reveals the framework is strongest where its steady-state assumption holds (Utilities 79%, Real Estate 67%, Staples 59%) and weakest where the assumption is most violated (Technology 43%).

### 9.3 What the framework *does not* claim

The Brina Gap is not a complete asset-pricing model. It does not produce expected returns, volatilities, or position sizes. It produces a *directional prediction*: firms in particular quadrants are more likely than chance to outperform or underperform the market over 5 years. Translating that into a portfolio construction is a separate problem.

The Brina Gap is not a high-frequency signal. The 5-year forward window is essential: the framework's mechanism is the eventual market recognition of the gap, which takes years to play out. The framework does not predict short-horizon returns.

The Brina Gap is not survivorship-bias-free in the strict sense: the S&P 500 itself is a selected universe (firms must be US-listed, large-cap, and meet liquidity/profitability requirements). Application to small-cap, international, or distressed equities is future work and explicitly not claimed here.

### 9.4 Implications for value investing

The empirical failure of single-axis MoS on this panel is, we believe, the paper's most consequential finding for practitioners. It is *not* that value investing fails — Buffett's track record alone falsifies any such claim — but that the *operational* form in which MoS is typically applied (price < DCF intrinsic value) is, on a broad equity universe, a noise signal. What works is the two-axis decomposition: firm economics *and* market expectations, considered jointly.

This is, we contend, the unspoken practice of every successful value investor. The Brina Gap names it.

### 9.5 Implications for academic finance

The Brina Gap is a *firm-specific*, *fundamentals-derived* signal that produces directionally significant excess returns versus a market benchmark. It is therefore a candidate factor in factor models, in spirit similar to Fama-French's value and quality factors but mechanically distinct: it is not a static cross-sectional sort but a year-by-year scalar computed per firm.

Whether the Brina Gap survives orthogonalization against the standard factors (market, size, value, profitability, investment) is empirical work we do not perform in full here. The tercile evidence of §8.5 indicates the Gap is not subsumed by profitability alone, but a complete factor-model test — Fama-French five-factor and Novy-Marx (2013) gross profitability regressions on a Gap-sorted long-short portfolio, with returns rather than directional hit rates — is reserved for a subsequent paper. The framework is offered here as a refinement of the conceptual case for value investing, not as a factor-model contribution.

---

## 10. Limitations

- **Universe**: US large-cap S&P 500 only. Applicability to small-cap, international, and emerging markets is untested.
- **Forward window**: 5-year horizon by design (the framework's mechanism is multi-year gap closure). Shorter horizons are not tested.

- **Sample window**: 2010–2024. Includes the post-GFC bull market, the COVID shock, and the 2022 rate-shock. Does not include the 1990s tech boom or 1970s stagflation; out-of-sample robustness in different macro regimes is future work.
- **Owner-earnings DCF**: our intrinsic-value methodology is one of several reasonable choices. A defender of a different MoS specification can substitute via the published pipeline.
- **Survivorship**: point-in-time S&P 500 membership mitigates but does not eliminate survivor bias inside the index.
- **Pending forward windows**: ~830 firm-years (FY2020+) have 5-year forward windows that complete after 2026-05-25. These are excluded from the primary headline and will be added in subsequent revisions as data matures. The maturing windows provide the out-of-sample panel on which §12's pre-registered refinements are to be tested.

---

## 11. Conclusion

The Brina Gap is a single, signed, falsifiable scalar that reframes the price-versus-value question by decomposing it into two independently estimable terms: the firm's fundamental growth ceiling ($g_f$ = ROIC × RR) and the market-implied growth rate from a reverse DCF ($g^*$). It is offered first as a theoretical construct — replacing the analyst-defined intrinsic value at the heart of the Margin of Safety with two externally verifiable quantities — and second as an empirical claim.

The empirical results, on the full S&P 500 panel with pre-registered universe and citation-backed exclusions, support the theoretical claim:

> The classical Margin of Safety predicts at 48.2% (p=0.94). The Brina Gap predicts at 52.6% (p=0.017) as a standalone signal and 55.1% (p=0.0003) on the strong-signal subset (|Gap| > 3pp, equivalently the four directional Matrix quadrants). On the firm-years where MoS and the Brina Gap disagree, the Brina Gap is correct 55% of the time, the Margin of Safety 45%. The Brina Gap is not a refinement of the Margin of Safety but a distinct measurement that outperforms it where the two disagree.

The framework's predictive content is asymmetric — strong as a negative screen (Value Trap 67%, Expensive Hype 59%) and weak as a positive screen on the broad universe. Per-sector heterogeneity reveals the framework operates best on firms with stable ROIC (Utilities 79%) and weakest on firms with rising or volatile ROIC (Technology 43%) — a scope condition exactly aligned with the framework's steady-state derivation.

The Brina Gap is, we contend, the quantitative form of what value investors have always sought: a measure of mispricing grounded in a firm's own economics rather than in the observer's assumptions. Unlike the Margin of Safety, it is falsifiable — and where the two disagree, it is the one more often proved right.

---

## 12. Refinements for v3.0 — Pre-Registered Hypotheses

The v2.0 results reveal three structural patterns that suggest refinements to the framework. These cannot be tested on the v2.0 panel without HARKing (Hypothesizing After Results are Known); they are pre-registered here as falsifiable hypotheses for v3.0, to be tested on (a) the FY2020+

forward-return windows of the v2.0 panel as they mature (windows complete 2025–2029) and/or (b) an out-of-sample international universe (e.g., S&P Europe 350, MSCI Japan large-cap).

This pre-registration is **dated 2026-05-25** and locked at v2.0 publication. v3.0 will report the pre-registered tests below and the results, including failures.

### 12.1 Refinement A — The ROIC-Trajectory Matrix

**Motivation**. §4.5 demonstrates that the MoS axis of the v2.0 Matrix contributes no incremental directional information beyond the |Gap|>3pp confidence band. §2.5 attributes this to the value-trap selection bias documented in Lakonishok-Shleifer-Vishny (1994), Sloan (1996), and Penman (2007): MoS+ is not an independent confirmation of cheapness, but a contaminated signal. A genuinely independent second axis must measure something the Gap cannot see, *and* must not carry MoS's bias.

**Proposed second axis: ROIC trajectory**, operationalized as the comparison between *marginal ROIC* ($\Delta E_{t+1}/\Delta IC_t$) and *average ROIC* ($E_t/IC_t$):

- Marginal ROIC > Average ROIC: expanding economics (the firm's reinvestment is producing higher returns than its historical capital base implies)
- Marginal ROIC ≤ Average ROIC: contracting economics

This is mechanically computable from the same panel inputs as $g_f$, requires no analyst assumption, and directly addresses Failure Mode 2 (Quality Compounder Blind Spot, §7.2).

**Proposed v3.0 Matrix**:

| | Gap < −3pp | Gap > +3pp |
|---|---|---|
| **Marginal ROIC > Average ROIC** (expanding economics) | Justified Premium — abstain | **Compounding Bargain** — strongest BUY |
| **Marginal ROIC ≤ Average ROIC** (contracting economics) | **Terminal Decline** — strongest SELL | Reinvestment Mirage — abstain |

**Pre-registered prediction A**: On the out-of-sample panel, the *Compounding Bargain* and *Terminal Decline* quadrants will jointly exceed a 60% directional hit rate against a 50% null (one-sided $p < 0.01$). The *Justified Premium* and *Reinvestment Mirage* quadrants are predicted to be non-significant at the 5% level (i.e., the framework correctly identifies these as abstain regions).

### 12.2 Refinement B — Asymmetric Application Rule

**Motivation**. §6.2 reveals that the v2.0 framework's positive predictions on the broad S&P 500 are anti-predictive (Underestimated Growth 48%, Double Discount 41%) while negative predictions are strong (Value Trap 67%, Expensive Hype 59%). The asymmetric assumption load identified in §6.2 — positive predictions require both ROIC stability *and* market recognition, negative require only correct identification of unsustainable implied growth — predicts this asymmetry.

**Pre-registered prediction B**: On the out-of-sample panel:

- Negative predictions (Gap < −3pp) on the broad universe will exceed 60% hit rate against a 50% null (one-sided $p < 0.01$).

- Positive predictions (Gap $>$ +3pp) on the broad universe will *not* exceed 55% hit rate (the v2.0 asymmetry will replicate).
- Positive predictions *restricted to stable-ROIC sectors* (Utilities, Real Estate, Consumer Staples, Healthcare — as identified in §6.3) will exceed 60% hit rate against a 50% null (one-sided $p < 0.01$).

### 12.3 Refinement C — Smoothed $g_f$ Input

**Motivation**. Single-year NOPAT and IC are noisy due to M&A, working-capital fluctuations, and one-time items. The v2.0 framework addresses this via boundary conditions and the cited-events database, but residual single-year noise still distorts $g_f$. Multi-year smoothing of ROIC reduces variance without introducing analyst input.

**Proposed input refinement**: Replace single-year ROIC in the $g_f$ identity with a 3-year trailing average:

$$g_f^{\text{smooth}} \equiv \overline{\text{ROIC}}_{t-2..t} \times b_t$$

**Pre-registered prediction C**: On the out-of-sample panel, $g_f^{\text{smooth}}$ will outperform single-year $g_f$ in pooled directional hit rate by at least 1 percentage point.

### 12.4 Refinement D — Tiered Confidence Bands

**Motivation**. §7.4 identifies that $|\text{Gap}| \in [0, 5\text{pp}]$ is below the framework's signal-to-noise floor. The v2.0 framework uses a single neutral band at 3pp; a graded tier structure should improve signal selectivity by stratifying predictions by signal strength.

**Proposed tier structure**:

- $|\text{Gap}| \in [3\text{pp}, 10\text{pp}]$: Medium confidence
- $|\text{Gap}| \in [10\text{pp}, 25\text{pp}]$: High confidence
- $|\text{Gap}| > 25\text{pp}$: Suspended (likely measurement artifact from boundary-condition-edge cases)

**Pre-registered prediction D**: On the out-of-sample panel, the High-confidence tier will exceed 60% directional hit rate and outperform the Medium-confidence tier by at least 3 percentage points.

### 12.5 What this section is and is not

This section is a *commitment device*. By stating these refinements as pre-registered predictions with specific thresholds, we forfeit the option of later cherry-picking which refinement to publish based on out-of-sample outcomes. v3.0 will report all four predictions A–D, including any that fail.

This section is *not* a claim about v2.0. The v2.0 results stand as published; nothing here modifies them. Refinements A–D are testable forward predictions, not retrospective adjustments. The companion `pre_registration_v3.md` in the published repository contains an identical statement of these hypotheses, dated and committed at v2.0 release.

---

## Appendix A — Worked Example: NVIDIA FY2012

NVIDIA's fiscal year 2012 ended January 29, 2012. At that point the company was a graphics-chip vendor pre-dating the data-center / AI cycle by approximately five years. We walk through the Brina Gap computation as it would have been visible to any reader of the 10-K filed in March 2012.

### A.1 Inputs (from the FY2012 10-K, SEC EDGAR)

| Quantity | Value | Source |
|---|---:|---|
| NOPAT | \$568M | EBIT × (1 − effective tax rate) |
| Invested Capital (avg) | \$1,782M | Average of beginning and ending IC |
| Reinvestment (net of D&A) | \$278M | Capex + ΔWC + Acquisitions − D&A |
| Free Cash Flow | \$290M | NOPAT − reinvestment |
| Year-end close | \$14.91 | NASDAQ |
| Shares outstanding | 612.2M | 10-K |
| Total debt | \$0 | 10-K |
| Cash & equivalents | \$668M | 10-K |

### A.2 Computed quantities

| Quantity | Computation | Value |
|---|---|---:|
| **ROIC** | \$568M / \$1,782M | **34.7%** |
| **Reinvestment Rate** $b$ | \$278M / \$568M | **49.0%** |
| **Fundamental growth ceiling** $g_f$ | 34.7% × 49.0% | **17.0%** |
| Market cap | 14.91 × 612.2M | \$9.13B |
| **Enterprise Value** | 9.13 + 0 − 0.67 | **\$8.46B** |
| EV / NOPAT | 8.46 / 0.568 | 14.9× |
| **Market-implied growth rate** $g^*$ | Reverse DCF on $EV$, NOPAT, $r$=10%, $g_T$=3%, $N$=10 | **3.4%** |
| **Brina Gap** | $g_f - g^* = 17.0\% - 3.4\%$ | **+13.6 pp** |
| **Quadrant** | $g_f$ high, $g^*$ low | **Underestimated Growth** |

### A.3 Interpretation as of January 2012

The framework saw a firm with operating economics that could sustain ~17% earnings growth from internal reinvestment alone, trading at a price that — under standard discounting — required only 3.4% growth to justify. The gap was large (+13.6 pp), the direction was clear (the firm's economics supported much more growth than the market was pricing), and the quadrant assignment was unambiguous (Underestimated Growth).

The framework's *a priori* prediction was: positive 5-year excess return versus SPY.

### A.4 Realized 5-year outcome (FY2012 close to FY2017 close)

| Quantity | Value |
|---|---|
| NVDA annualized 5-year return | **+49.6 % / yr** |
| SPY annualized matched-window return | +11.7 % / yr |
| **Realized 5-year annualized excess return** | **+37.9 percentage points per year** |

NVIDIA outperformed SPY by 37.9 percentage points per year for the five years following the FY2012 fiscal close. The framework's prediction was directionally correct, and the magnitude of outperformance is the largest single-firm hit in the panel.

### A.5 What the framework did not predict

The framework predicted *direction* (positive excess return). It did not predict *magnitude* (37.9 pp/yr is far beyond what 13.6 pp of Gap suggests for a five-year window). The mechanism of the magnitude — the data-center and AI cycle that re-rated NVIDIA's ROIC upward by an order of magnitude — is exactly Failure Mode 2 (Quality Compounder Blind Spot, §7.2): the framework's steady-state assumption *under*-predicts firms whose ROIC subsequently rises. NVIDIA's $g_f$ of 17% was a *ceiling on the steady-state*; the firm broke through the ceiling by acquiring a new addressable market (data center). The framework was conservatively right in direction, missed the magnitude, and the miss is in the favorable direction.

This is, we contend, the most useful kind of failure to have in a value-investing framework.

## Appendix B — Replication

All code, the 6,428-row dataset, the 228-event corporate-events database, the manual review queue, and per-quadrant bootstrap CIs are published at the companion repository:

- `pre_registration.md` — pre-registration with two dated amendments
- `outputs/file_01_full_dataset.csv` — the full panel
- `outputs/ticker_events_database.csv` — citation-backed exclusion log
- `outputs/statistical_results.md` — auto-generated results file
- `outputs/file_05_replication_guide.md` — methodology details
- `src/` — pipeline implementation (Python 3.12)

Any reader can reproduce every number in this paper from raw SEC filings using the published pipeline. The user-agent header for SEC EDGAR API access must comply with SEC's terms of use.

## Appendix C — Notation

| Symbol | Meaning |
|---|---|
| $E_t$ | NOPAT at time $t$ |
| $IC_t$ | Invested capital at time $t$ (average of beginning-of-year and end-of-year) |
| $\text{ROIC}_t$ | $E_t/IC_t$ |

| Symbol | Meaning |
|---|---|
| $b$ | Reinvestment rate (plowback fraction) |
| $g_f$ | Fundamental growth ceiling = ROIC $\times$ $b$ |
| $g^*$ | Market-implied growth rate from reverse DCF on EV |
| $r$ | Discount rate (uniform 10%) |
| $g_T$ | Terminal growth rate (3% standardized) |
| $FCF_t$ | Free cash flow to the firm at time $t$ |
| $EV_t$ | Enterprise value at $t$ |
| Gap | $g_f - g^*$ |
| MoS | Margin of Safety = (intrinsic value − market cap) / market cap |
| BC1–BC6 | Boundary conditions (suspension criteria) |

## Declarations

**Data and code availability.** The full panel dataset (6,428 ticker-years), the corporate-events database, the pre-registration documents, and the complete replication pipeline are published as companion materials (Appendix B; Zenodo and the project repository). Every number in this paper is reproducible from public SEC filings using the published code.

**Use of AI tools.** The author developed the Brina Gap framework, designed the research methodology, collected and analyzed all data, and drew all conclusions independently. Claude (Anthropic) was used solely as a writing assistant to improve the readability and formatting of the prose.

**Competing interests.** The author is the founder of Zyberno.com, a venture that applies the Brina Gap framework, and therefore has a professional and prospective financial interest in the framework — though he presently draws no income from the venture and funds its operation himself. This is disclosed in the interest of full transparency. The study is deliberately designed to be resistant to author bias: the universe and methodology were pre-registered before any results were computed (with dated amendments); every firm-year exclusion is documented with a verifiable citation; the framework's failure modes, weakest sectors, and discount-rate sensitivity are reported in full rather than suppressed; and the complete dataset and replication pipeline are published so that any reader can independently reproduce — or refute — every result in this paper.

**Funding.** This research received no external funding.

---

## References

Bessembinder, H. (2018). *Do stocks outperform Treasury bills?* Journal of Financial Economics 129(3), 440–461.

Buffett, W. (1986). *Letter to shareholders of Berkshire Hathaway Inc.* (introducing the owner-earnings concept).

Damodaran, A. (2002, and continuing editions). *Investment Valuation: Tools and Techniques for Determining the Value of Any Asset.* Wiley.

Fama, E. F., and French, K. R. (1993). *Common risk factors in the returns on stocks and bonds*. Journal of Financial Economics 33(1), 3–56.

Graham, B., and Dodd, D. L. (1934). *Security Analysis*. McGraw-Hill.

Graham, B. (1949). *The Intelligent Investor*. Harper & Brothers.

Greenblatt, J. (2005). *The Little Book That Beats the Market*. Wiley.

Greenwald, B., Kahn, J., Sonkin, P., and van Biema, M. (2004). *Value Investing: From Graham to Buffett and Beyond*. Wiley.

Koller, T., Goedhart, M., and Wessels, D. (2020). *Valuation: Measuring and Managing the Value of Companies* (7th ed.). McKinsey & Company / Wiley.

Lakonishok, J., Shleifer, A., and Vishny, R. W. (1994). *Contrarian investment, extrapolation, and risk*. Journal of Finance 49(5), 1541–1578.

Mauboussin, M. J., and Rappaport, A. (2001). *Expectations Investing: Reading Stock Prices for Better Returns*. Harvard Business School Press.

Novy-Marx, R. (2013). *The other side of value: The gross profitability premium*. Journal of Financial Economics 108(1), 1–28.

Penman, S. H. (2007). *Financial Statement Analysis and Security Valuation*. McGraw-Hill (3rd ed.).

Penman, S. H. (2010). *Accounting for Value*. Columbia University Press.

Ross, S. A., Westerfield, R. W., and Jaffe, J. F. (various editions). *Corporate Finance*. McGraw-Hill. (See Chapter 5 on the sustainable growth identity.)

Sloan, R. G. (1996). *Do stock prices fully reflect information in accruals and cash flows about future earnings?* The Accounting Review 71(3), 289–315.

---

*Working Paper v2.0 — comments to fabio@fabiobrina.com. Companion materials and replication pipeline published at Zenodo and the Zyberno Research repository. Intended for working-paper circulation; not yet submitted to a journal.*