

Market Prism: Narrative Exhaustion, Valuation Disequilibrium, and State-Dependent Equity Repricing

A point-in-time event-study framework for distinguishing narrative failure from narrative recovery

Tara Walsh

Market Prism · marketprism.co · tara@vtlbranding.com

Working Paper — September 2026 (draft; not peer reviewed)

Abstract

This paper studies narrative exhaustion as a state transition rather than a directional trading signal. Using Market Prism’s proprietary point-in-time equity panel, we examine high-confidence exhaustion episodes from the live measurement period beginning March 2, 2026 and condition subsequent returns on fair-value divergence, narrative verdict transitions, narrative-energy trajectories, and thematic context. After de-duplicating consecutive exhausted observations into distinct episodes, deeply undervalued securities (fair-value divergence below -20%) exhibit positive subsequent drift: across 196 episodes, mean returns are $+1.74\%$ at five trading days, $+3.04\%$ at ten days, and $+2.54\%$ at twenty days; the twenty-day median is $+2.88\%$ and 60.4% of observations are positive. Marginal undervaluation does not display the same behavior. On the overvalued side, exhaustion alone is not bearish. A narrower state characterized by persistent narrative risk and collapsing narrative energy produces a -1.39% mean and -1.83% median twenty-day return across 58 de-duplicated episodes; the SPY-adjusted mean and median are -1.53% and -2.66% , respectively. In the later live measurement period, the same state records a -3.90% mean and -3.64% median twenty-day return while other overvalued exhaustion states remain positive. Among undervalued securities in the later measurement period, transitions from an exhausted narrative toward renewed structural support or accumulation are associated with $+6.59\%$ mean, $+5.42\%$ median, and $+5.39\%$ SPY-relative twenty-day returns ($n = 25$), though this subgroup is exploratory. The evidence supports a conditional interpretation: exhaustion marks depletion of an informational state, while subsequent direction depends on valuation displacement and whether narrative structure deteriorates, persists, or reconstructs.

Keywords: narrative economics; narrative exhaustion; fair-value divergence; structural mispricing; state transitions; information decay; event studies; point-in-time panel data.

JEL classification: G11, G12, G14, G17, C58.

1. Introduction

Market narratives are commonly measured as levels: sentiment is positive or negative, attention is high or low, and textual tone is compressed into a contemporaneous score. Such representations omit a dynamic question that is central to price formation: what happens when the information process sustaining a narrative is depleted?

Market Prism treats narrative exhaustion as a transition state. A narrative may lose energy while the security remains fundamentally supported; it may remain risky while informational support collapses; or an exhausted narrative may reconstruct into a new state of structural support. These paths need not imply the same subsequent price direction. The central empirical question of this paper is therefore not whether

exhaustion predicts reversal, but whether the interaction between exhaustion, valuation displacement, and narrative-state evolution separates subsequent return distributions.

This paper extends the state-space framework developed in prior Market Prism work, which fixes a non-reflexive fundamental reference frame, measures price displacement from that frame, and models narrative intensity and decay dynamically. In that framework, fair-value displacement is a measurement rather than a forecast, and narrative intensity is interpreted jointly with decay and position in the information cycle. The present paper isolates one implication of that architecture: exhaustion should have different price consequences depending on where price sits relative to the fundamental frame and on what happens to narrative structure around the exhaustion boundary.

The paper makes four contributions:

- (1) **Undervalued exhaustion is heterogeneous.** Deep discounts are followed by positive drift, while marginal discounts are not (Section 5.1).
- (2) **Generic overvaluation is not a short signal.** Generic overvalued exhaustion is not a reliable short condition (Section 5.2).
- (3) **Persistent risk plus collapsing energy.** A narrower overvalued state — persistent narrative risk combined with collapsing narrative energy — has a materially weaker subsequent return distribution (Sections 5.3–5.4).
- (4) **Undervalued narrative recovery.** Exploratory evidence identifies a mirrored recovery state in which undervalued exhausted securities transition toward renewed structural support and subsequently reprice upward (Section 5.5).

The central proposition is state-dependent rather than directional: narrative exhaustion marks depletion of an informational state, while subsequent price direction depends jointly on valuation displacement and the structure of what follows.

2. The State-Transition Framework

2.1 Valuation displacement

Let p_t denote the observable spot price of a security at time t and $V_{ground,t}$ the point-in-time fundamental reference frame. Define normalized displacement as:

$$\Delta_t = (p_t - V_{ground,t}) / V_{ground,t}.$$

Negative displacement denotes price below the modeled fundamental frame; positive displacement denotes price above it. In production this concept is operationalized through fair-value divergence (FVD). As in the earlier Market Prism framework, displacement is a state coordinate rather than a directional forecast. A security may remain displaced for an extended period if narrative support persists.

2.2 Narrative exhaustion

Narrative exhaustion denotes depletion of the informational process sustaining the dominant story. It is deliberately separated from price direction. A bullish narrative can exhaust before price falls; a bearish narrative can exhaust while price is deeply discounted; and a depleted narrative can subsequently rebuild.

The relevant object is therefore a joint state rather than a single label. In the notation of this paper, the state contains at minimum valuation displacement, exhaustion status, narrative verdict, and narrative-energy trajectory. The empirical question is whether particular combinations of these coordinates are followed by distinct return distributions.

2.3 Narrative Failure and Narrative Recovery

We distinguish two candidate transition mechanisms. *Narrative Failure* occurs when an overvalued exhausted security remains in a risk-class verdict while narrative energy collapses. The intuition is not that overvaluation must reverse, but that an already-stretched price becomes vulnerable when the informational structure supporting it deteriorates.

Narrative Recovery occurs when an undervalued exhausted security migrates from an exhausted verdict toward renewed structural support or accumulation. Here the depletion of a prior narrative is followed by reconstruction while price remains below the modeled fundamental frame. The framework therefore predicts opposite subsequent distributions for these two states without claiming that exhaustion alone has a directional sign.

3. Data and Point-in-Time Discipline

The study uses the proprietary Market Prism production panel and is restricted to the live point-in-time period beginning March 2, 2026. High-confidence exhaustion observations require status = EXHAUSTED and exhaustion confidence of at least 0.80. Valuation state is measured using the point-in-time fair-value divergence available at the event timestamp.

Forward prices are taken from the split-adjusted gold daily-bar series. Benchmark returns are reconstructed directly from the SPY daily-bar series rather than from precomputed Market Prism SPY statistics. Sampled SPY closes were independently checked before use because earlier internal benchmark fields had known inconsistencies.

A central design choice is event de-duplication. Consecutive exhausted ticker-days are not treated as independent observations. The principal event sample requires a break of more than five calendar days before a new exhaustion episode for the same ticker on the same side of modeled fair value. This materially reduces the nominal sample size and is intentionally conservative: the purpose is to study distinct exhaustion episodes rather than repeated daily measurements of the same episode.

The paper also distinguishes the early live measurement period (March–June 2026) from the later live measurement period (July–September 2026). This is a measurement-era comparison, not a clean claim of prospective out-of-sample validation. Market Prism’s scoring architecture improved during the live period, so later results may partly reflect improved state measurement.

4. Empirical Design

Returns are measured from the event-date close to the close five, ten, and twenty trading sessions later. We report arithmetic means, medians, directional frequencies, trimmed means where relevant, and SPY-relative returns. Statistics at each horizon use only episodes whose forward window is complete at the data cut-off, so the effective twenty-day sample is smaller than the episode count; table notes report it.

For overvalued exhaustion, the principal deterioration state is defined as a risk-class verdict five sessions before exhaustion, a risk-class verdict at exhaustion, and a decline of at least 15 points in narrative energy over that interval. The risk class groups verdicts whose semantics indicate narrative weakness, trap risk, or downside pressure. For undervalued exhaustion, we examine both valuation-severity buckets and recovery transitions from Exhausted Narrative toward Structurally Supported or Accumulate.

The research process is explicitly sequential. Broad valuation-bucket results were discovered first; narrower verdict-energy and theme combinations were then investigated. For that reason, subgroup findings are labeled exploratory rather than confirmatory. Exact proprietary estimators, field weights, and production thresholds beyond those necessary to define the reported tests remain undisclosed.

5. Empirical Results

5.1 Deep undervaluation at exhaustion

FVD at exhaustion	N	5D mean	5D median	10D mean	20D mean	20D median	20D positive
< -20%	196	+1.74%	+0.55%	+3.04%	+2.54%	+2.88%	60.4%
-20% to -10%	95	+0.58%	+1.34%	+0.53%	+0.42%	+0.47%	51.9%
-10% to -5%	53	+0.26%	+0.61%	+1.62%	+0.43%	-0.27%	47.7%
-5% to 0%	108	-0.17%	+0.05%	-0.93%	+0.55%	-1.58%	43.7%

Table 1. De-duplicated undervalued exhaustion episodes by valuation displacement. The broad result is threshold-like rather than perfectly monotonic. N counts episodes with an event-date close; twenty-day statistics use episodes with a completed twenty-session window ($n = 159, 77, 44$ and 87 , respectively).

The deepest discount bucket is the strongest broad finding. Across 196 de-duplicated episodes with FVD below -20%, the five-day mean return is +1.74%, the ten-day mean is +3.04%, and the twenty-day mean is +2.54%. The twenty-day median is +2.88%, and 60.4% of observations are positive.

Marginal undervaluation behaves differently. Securities between -5% and 0% FVD have a -0.93% mean ten-day return, a -1.58% median twenty-day return, and only 43.7% positive observations at twenty days. The evidence therefore does not support the generic statement that an exhausted stock below fair value should rise. The effect is concentrated in materially discounted states.

5.2 Generic overvaluation is not a short signal

After event de-duplication, generic overvalued exhaustion remains positive on average over several horizons. This rejects the simplest symmetric mean-reversion hypothesis. Overvaluation appears to create vulnerability, not a deterministic short signal. The path into exhaustion is required to identify the weaker subset.

5.3 Persistent risk plus collapsing energy

Statistic	20D result
Episodes	58
Episodes with completed 20D window	52
Distinct tickers	31
Mean return	−1.39%
Median return	−1.83%
5–95% trimmed mean	−1.90%
Negative observations	59.6%
SPY-adjusted mean	−1.53%
SPY-adjusted median	−2.66%
Naive mean <i>t</i> -statistic	−0.75
Ticker-clustered <i>t</i> -statistic (preliminary)	−0.78

Table 2. De-duplicated overvalued Risk-to-Risk + Energy Collapse episodes. Robust location measures are consistently negative, but conventional mean-based statistical significance is not yet established. Twenty-day statistics use the 52 episodes (30 tickers) with a completed twenty-session window; the clustered statistic applies the $G/(G - 1)$ small-sample correction.

The mean, median, trimmed mean, and SPY-relative median all point in the same direction. Trimming the 5% tails makes the mean more negative, indicating that the result is not created by a small number of extreme losses. The SPY-relative median is −2.66%, compared with a raw median of −1.83%.

The limitation is dispersion. Individual-equity twenty-day returns are noisy, and the sample standard deviation is approximately 13.3%. The present sample therefore does not support a strong significance claim from the arithmetic mean. The result is economically suggestive and distributionally coherent, but inferentially imprecise.

5.4 Measurement-era discrimination

Measurement era	Risk → Risk + Collapse	Other overvalued exhaustion	Approx. separation
Mar–Jun	−0.06%	+0.12%	−0.18 pp
Jul–Sep	−3.90%	+2.56%	−6.46 pp

Table 3. Mean twenty-day returns by measurement era. The difference-in-differences estimate is approximately −6.28 percentage points, with a preliminary ticker-clustered *t*-statistic of approximately −1.82 (−1.57 under independent cell variances). Cells contain 34 and 212 (Mar–Jun) and 18 and 249 (Jul–Sep) episodes with completed twenty-day windows.

The later period exhibits substantially greater state separation. The deterioration state records a −3.90% mean and −3.64% median twenty-day return, while other overvalued exhaustion states record a +2.56% mean and +1.32% median. Equal-weighting tickers produces a similar later contrast (−3.67% versus +2.08%).

The interpretation is measurement-based rather than purely predictive. Market Prism’s scoring architecture improved during the live period. The widening separation is therefore consistent with improved classification of latent narrative states, but it cannot be described as a clean frozen-model out-of-sample result.

5.5 Undervalued narrative recovery

Later-era state	N	5D mean	10D mean	20D mean	20D median	20D SPY-adj usted
Other undervalued exhaustion	181	+0.88%	+0.75%	+1.38%	+1.48%	+0.13%
Support → Support + Energy Rebuild	64	+0.86%	+1.12%	+3.29%	+1.15%	+1.67%
Exhausted → Support/Accumulate	25	+3.85%	+3.51%	+6.59%	+5.42%	+5.39%

Table 4. Later measurement-era undervalued states. The recovery-transition subgroup is exploratory because $n = 25$, but both raw and SPY-relative returns are economically large. Twenty-day statistics use episodes with a completed twenty-session window ($n = 121, 45$ and 19 , respectively).

The recovery transition is the clearest candidate mechanism on the undervalued side. A security begins from an exhausted narrative state while priced below modeled fair value, then migrates toward renewed structural support or accumulation. The subsequent twenty-day mean is +6.59%, the median is +5.42%, and the SPY-relative mean is +5.39%.

A broader reinforcement state — Support to Support with rebuilding energy — contains 64 episodes and records a +3.29% mean twenty-day return and +1.67% SPY-relative mean. This larger sample suggests a slower positive repricing process, although it is not as economically large as the recovery-transition subgroup.

5.6 Thematic and verdict-path evidence

Theme-conditioned cells are smaller and are therefore treated as hypothesis-generating. Within overvalued exhaustion, persistent Company-Specific narratives combined with persistent risk and collapsing energy produced a -4.37% median twenty-day return with 70.6% of observations negative ($n = 20$ episodes, 17 with completed twenty-day windows). A Company-Specific Support-to-Support + Energy Collapse subgroup produced a -4.89% twenty-day median with 77.8% negative ($n = 14$ episodes, 9 with completed windows).

These cells are too small for headline claims, but they suggest that energy depletion may matter even when the nominal verdict remains supportive. The thematic results motivate prospective tests of whether company-specific narratives unwind differently from earnings, regulatory, geopolitical, or sector narratives.

6. Interpretation: Exhaustion as a State Transition

The empirical evidence rejects the simplest interpretation of narrative exhaustion. Exhaustion is neither intrinsically bullish nor intrinsically bearish. It marks depletion of an informational state.

On the overvalued side, valuation displacement appears to create vulnerability, while narrative deterioration determines whether that vulnerability is realized. Persistent risk combined with collapsing energy is associated with a weaker twenty-day distribution, especially in the later measurement era. Generic overvalued exhaustion, by contrast, is not a stable short condition.

On the undervalued side, the broadest evidence is strongest when the discount is substantial. Deeply undervalued exhausted securities subsequently drift upward, while marginally undervalued exhausted securities do not. A narrower recovery transition — from Exhausted Narrative toward renewed support or accumulation — is associated with considerably larger positive repricing.

The framework can therefore be described as a two-sided state machine. Narrative Failure is the transition from stretched valuation plus depleted risky narrative toward negative repricing. Narrative Recovery is the transition from discounted valuation plus depleted narrative toward renewed structural support and positive repricing. In this formulation, valuation is potential and narrative evolution governs the release.

7. Robustness, Inference, and Boundary Conditions

Event de-duplication is central. Treating consecutive exhausted days as independent observations materially inflates sample size and understates dependence. The episode-level results therefore require a spacing rule before a ticker can generate a new event.

Benchmark adjustment is reconstructed from the checked daily SPY series rather than inherited from precomputed benchmark statistics. On the principal overvalued deterioration state, SPY adjustment makes the median more negative (-2.66% versus -1.83%), which argues against the effect being solely a broad-market artifact.

Robust location measures support the sign of the overvalued deterioration effect, but conventional mean-based significance is not yet established. The preliminary ticker-clustered t -statistic is approximately -0.78 for the 52 of 58 episodes with completed twenty-day windows. The measurement-era interaction is economically large but has a preliminary ticker-clustered t -statistic of approximately -1.82 (-1.57 under independent cell variances). We therefore report effect sizes and uncertainty rather than promoting the result as statistically established alpha.

The later measurement era should not be labeled fully out-of-sample. Market Prism scoring improved during the live period, so later results may reflect improved measurement. This is itself an empirical observation — state separation widened — but it prevents a clean prospective interpretation.

The strongest broad undervalued result, deep discount at exhaustion, is descriptive and based on a materially larger sample. The narrower recovery-transition result is economically large but exploratory. Prospective accumulation of additional events under a frozen scoring architecture is required before stronger inferential claims are appropriate.

8. Relation to the Literature

This paper sits at the intersection of narrative economics, media and attention effects, behavioral mispricing, and empirical evaluation. Shiller (2017, 2019) motivates narratives as economic forces; Tetlock (2007) and Da, Engelberg and Gao (2011) establish that media tone and attention contain market information. Market Prism contributes a dynamic state representation in which narrative intensity can deplete and reconstruct, conditioned on a non-reflexive valuation coordinate.

The present paper is also a direct extension of the Market Prism state-space framework described in Walsh (2026), which formalizes a fundamental reference frame, fair-value displacement, narrative decay, and a point-in-time measurement architecture. The contribution here is narrower: it asks what happens at the exhaustion boundary and shows that direction depends on the state surrounding that boundary rather than exhaustion itself.

The methodological emphasis follows the evaluation literature's warnings about overlapping observations, multiple testing, and backtest overfitting (Harvey, Liu and Zhu, 2016; Bailey et al., 2014;

López de Prado, 2018). The present sample is intentionally framed conservatively: de-duplicated episodes replace ticker-days as the principal unit, medians and trimmed means accompany arithmetic means, and discovered subgroups are not promoted to confirmatory claims.

9. Limitations and Boundary Conditions

Short live history. Dense live point-in-time evidence begins in 2026. Several exhaustion states are episodic, and some subgroup counts remain small.

Changing measurement architecture. The scoring system improved during the live period. This increases the value of later measurements but complicates claims of strict prospective validation. The early-versus-later comparison is therefore a measurement-era analysis rather than a frozen-model OOS test.

Fair-value model risk. Fair value is a model-based reference frame, not an observable truth. Errors in the anchor can misclassify apparent over- or undervaluation. The strongest claims concern conditional return distributions given the system's measured state, not proof that an asset was objectively mispriced.

Multiple-testing risk. Exact theme-verdict-energy combinations were explored after broader patterns were observed. Those cells are treated as hypothesis-generating and require prospective replication.

No causal identification. The analysis documents conditional associations. It does not establish that narrative depletion itself causes subsequent repricing or exclude omitted common factors.

Not investment advice. All statistics are historical and computed on a proprietary research panel. Nothing in this paper is a recommendation to transact in any security.

10. Conclusion

Narrative exhaustion is better understood as a boundary between informational states than as a directional signal. Deep undervaluation at exhaustion is followed by positive drift in the present live sample, while marginal undervaluation is not. Generic overvalued exhaustion does not reliably reverse; a narrower deterioration state characterized by persistent risk and collapsing narrative energy has a weaker subsequent distribution, especially under the later measurement architecture. Conversely, undervalued securities that transition from exhaustion toward renewed structural support display strong positive repricing in the later measurement period.

The empirical picture is therefore conditional rather than symmetric. Valuation displacement determines vulnerability or recovery potential; exhaustion marks depletion; and narrative evolution determines whether the depleted state resolves through failure, reconstruction, or continuation. The principal open task is prospective validation under a frozen scoring architecture. Until then, the evidence supports a state-transition taxonomy and several economically meaningful hypotheses, but not a claim of universal reversal or fully established standalone alpha.

Disclosures. The author is the sole architect of the system described. Market Prism and certain of its components are the subject of pending United States patent applications. This paper reports historical statistics computed on a proprietary point-in-time panel; underlying estimators, weights, and implementation are withheld as trade secrets. Field names and constructs are described at a level sufficient for evaluation but insufficient for replication. Nothing herein is investment advice. The author reports no external funding for this paper.

References

- Bailey, D. H., Borwein, J. M., López de Prado, M., and Zhu, Q. J. (2014). Pseudo-mathematics and financial charlatanism: the effects of backtest overfitting on out-of-sample performance. *Notices of the AMS*, 61(5), 458–471.
- Bailey, D. H., and López de Prado, M. (2014). The deflated Sharpe ratio: correcting for selection bias, backtest overfitting, and non-normality. *Journal of Portfolio Management*, 40(5), 94–107.
- Da, Z., Engelberg, J., and Gao, P. (2011). In search of attention. *Journal of Finance*, 66(5), 1461–1499.
- Harvey, C. R., Liu, Y., and Zhu, H. (2016). ... and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5–68.
- López de Prado, M. (2018). *Advances in Financial Machine Learning*. Wiley.
- Shiller, R. J. (2017). Narrative economics. *American Economic Review*, 107(4), 967–1004.
- Shiller, R. J. (2019). *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*. Princeton University Press.
- Tetlock, P. C. (2007). Giving content to investor sentiment: the role of media in the stock market. *Journal of Finance*, 62(3), 1139–1168.
- Walsh, T. (2026). Market Prism: A State-Space Framework for Structural Mispricing, Narrative Decay, and Sector Energy Reallocation in Equity Markets. Working paper.

Appendix A. Operational Definitions

Construct	Operational meaning in this paper
Fair-value divergence (FVD)	Point-in-time percentage displacement of price from the Market Prism fundamental reference frame.
Narrative exhaustion	High-confidence EXHAUSTED state in the production scorecard.
Risk persistence	Risk-class verdict five sessions before exhaustion and at the exhaustion event.
Energy collapse	Narrative-energy decline of at least 15 points over the five-session lead-in.
Narrative recovery	Transition from Exhausted Narrative toward Structurally Supported or Accumulate while FVD is negative.
Distinct episode	An exhaustion observation preceded by more than five calendar days without another qualifying exhaustion observation for that ticker on the same side of modeled fair value.

Table A1. Operational definitions used in the event study.

Appendix B. Claims Hierarchy

Broad descriptive claim. Deeply undervalued exhausted securities in the live PIT sample exhibit positive subsequent drift, with the strongest broad evidence in the $FVD < -20\%$ bucket.

Conditional descriptive claim. Generic overvalued exhaustion is not bearish; weaker subsequent returns are concentrated in states that combine overvaluation, persistent narrative risk, and collapsing narrative energy.

Exploratory mechanism claim. Undervalued exhaustion followed by renewed structural support or accumulation is associated with materially positive repricing in the later measurement era.

Not established. The present evidence does not establish causal effects, universal reversal, fully prospective OOS performance, or standalone alpha after all possible risk-factor adjustments.