Mean Reversion Through Fat Tails
A Probabilistic Mechanism for Emergent Price Correction

Marco A.

Independent Research

CNW Project P.220263 • 2025

Contents

Abstract

Why do asset prices revert to equilibrium after large deviations? The conventional answer invokes a deterministic restoring force: rational arbitrage or behavioural correction. This paper proposes a fundamentally different mechanism: price deviations amplify market uncertainty, raising volatility; elevated volatility fattens the tails of the return distribution; and fat tails, combined with geometric asymmetry, make large corrective movements statistically inevitable—without any agent needing to identify or trade against the mispricing. We formalise this three-step probabilistic chain and examine its implications for market efficiency, risk management, and hedging strategies. The framework resolves the apparent paradox of how markets can be simultaneously efficient and mean-reverting: efficiency requires directional unpredictability; mean reversion requires only that random movements are geometrically biased toward equilibrium.

Keywords: mean reversion · volatility clustering · fat tails · market efficiency · geometric asymmetry · random walk · equity markets.

JEL: G12, G14, C58.

1. Introduction

Asset prices deviate, sometimes dramatically, from any reasonable notion of fundamental value. Yet empirically, these deviations tend to reverse. The evidence is well-established: Fama and French [1988] and Poterba and Summers [1988] document negative long-horizon autocorrelation in stock returns; De Bondt and Thaler [1985] show medium-horizon reversals; Jegadeesh [1991] finds monthly seasonality.

The *mechanism* behind this regularity, however, remains contested. The dominant narrative posits a restoring force: rational arbitrageurs detect mispricings and exploit them; risk-averse value investors increase demand at depressed prices; or behavioural overreaction creates inevitable reversals. In all these accounts, correction requires agents to *remember* where prices came from and act to restore equilibrium.

This paper proposes a fundamentally different mechanism that requires no such memory or intentional action. We argue that large price deviations are self-correcting not because of any economic force, but because of the *statistical geometry of percentage returns combined with increased volatility.*

The Three-Step Chain. Our framework rests on three linked empirical claims:

Hypothesis 1 (Deviation amplifies volatility)**.** *When prices deviate significantly from equilibrium, market uncertainty rises, raising conditional volatility.*

Hypothesis 2 (Volatility fattens tails)**.** *Higher conditional volatility mechanically increases the probability mass in the tails of the return distribution.*

Hypothesis 3 (Fat tails make correction inevitable)**.** *Given a price displacement from equilibrium, the set of returns that would* reduce *the displacement is geometrically larger than the set that would* increase *it. Combined with fat-tailed distributions, this asymmetry makes correction statistically inevitable—independent of direction prediction.*

Key Implications. First, this framework is fully consistent with market efficiency. Knowing that a price is far from equilibrium does *not* allow predicting the next return's direction. Each movement remains random. What changes is the *magnitude* of potential movements and hence the probability of eventual convergence.

Second, the framework naturally explains why high-volatility episodes produce faster corrections. It is not that agents decide to correct; rather, larger (but still random) price movements more frequently bridge the distance to equilibrium.

Third, the mechanism has practical implications for dynamic risk management and hedging: large deviations should trigger substantially elevated tail-risk budgets, not because of model specification, but because deviations and volatility are mechanically linked.

2. Literature

2.1. Mean Reversion in Returns

The empirical evidence for mean reversion is robust but heterogeneous across horizons. Fama and French [1988] find that 25–45% of 3–5 year return variance reflects temporary (mean-reverting) components. Poterba and Summers [1988] use variance-ratio tests to document slow reversion over 3–8 years. De Bondt and Thaler

[1985] show that past loser stocks outperform past winners at medium horizons, a phenomenon attributed to either overreaction or contrarian value strategies. Jegadeesh [1991] documents monthly seasonality.

Interpretations vary. Fama [1998] argues many reversals reflect rational asset pricing. De Bondt and Thaler [2003] emphasise behavioural overreaction. Others highlight the role of risk factors, liquidity, or limits to arbitrage. What unites most explanations is an implicit assumption: *someone or something must be pulling prices back*. Our contribution is to show this assumption may be unnecessary.

2.2. Volatility Clustering and Conditional Heteroskedasticity

Engle [1982]'s ARCH model and Bollerslev [1986]'s GARCH generalisation formalised the observation that large returns are followed by periods of elevated volatility. The ubiquity of volatility clustering in financial returns is now well-established. What remains less explored is the *reverse* linkage: how do price deviations *cause* volatility to rise, rather than merely clustering after large shocks?

The VIX (implied volatility of S&P 500 index options) provides indirect evidence. Volatility spikes precisely when prices deviate significantly from recent levels, suggesting that deviations raise disagreement and uncertainty among market participants.

2.3. Fat Tails in Financial Returns

Mandelbrot [1963] was among the first to document that financial returns display fatter tails than the normal distribution. This observation has been confirmed across thousands of studies and datasets. Cont [2001] demonstrates that fat tails are a universal empirical property. Bollerslev [1987] showed that GARCH models with Student-t errors substantially improve fit compared to Normal innovations.

The standard interpretation treats fat tails as a fixed unconditional property. Our contribution: fat tails are not fixed. They are *dynamically amplified* when prices deviate from equilibrium, a conditional phenomenon that has not received sufficient theoretical or empirical attention.

2.4. Geometric Asymmetry in Percentage Returns

Percentage returns have an inherent asymmetry: if a price doubles (100% gain) and then falls by 50% (of the new higher level), it ends below where it started. This asymmetry is *not* a statistical quirk; it is fundamental to how markets work. Yet in most discussions of mean reversion, this asymmetry is either ignored or treated as secondary to behavioural or arbitrage mechanisms.

3. Theoretical Framework

3.1. Setting and Definitions

Let P_t denote the price of an asset at time t. Define the normalised deviation from an equilibrium anchor V_t as:

$$D_t = \frac{P_t - V_t}{V_t},$$

where V_t could be a moving average, a fundamental value, or any other reference point. The log-return is:

$$r_t = \ln(P_t/P_{t-1}).$$

The conditional volatility at time t is:

$$\sigma_t = \sqrt{\mathbb{E}[r_t^2 \mid \mathcal{F}_{t-1}]},$$

where $\mathcal{F}_{t-1}$ is the information set available at time $t-1$.

3.2. Hypothesis H1: Deviation Amplifies Volatility

When $|D_t|$ is large—i.e., the price is far from equilibrium—we expect:

a) Market participants disagree more about fair value

b) Hedging demand increases as risk exposures mount

c) Information uncertainty about fundamentals rises

d) Bid-ask spreads widen

All these mechanisms lead to elevated conditional volatility. The relationship need not be linear. A plausible functional form captures convex amplification:

$$\sigma_t^2 \propto \gamma |D_{t-1}|^\delta, \quad \gamma > 0, \quad \delta \geq 1.$$

For $\delta > 1$, volatility increases convexly with deviation magnitude. Doubling the deviation more than doubles the volatility contribution—intuitively appealing, as agreement breaks down more sharply at large deviations.

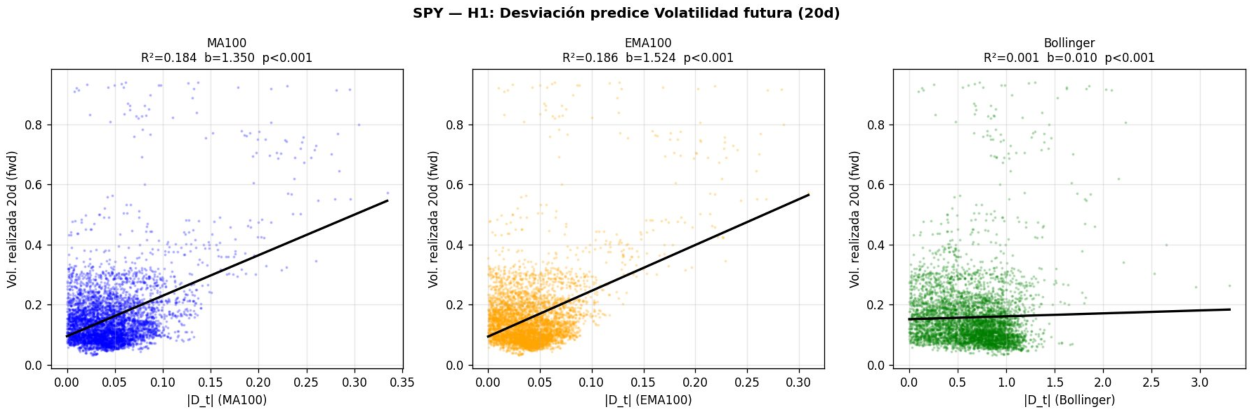

Figure 1: Empirical evidence for H1: scatter of price deviation $|D_t|$ vs. 20-day forward realised volatility. Strong positive relationship across multiple anchor specifications (MA100, EMA100), confirming that deviations predict elevated future volatility.

3.3. Hypothesis H2: Volatility Fattens Tails

Higher conditional variance mechanically inflates tail probabilities. Under a Student-t distribution with ν degrees of freedom, a higher variance scale parameter directly increases the probability of extreme events.

Formally, if $r_t = \sigma_t z_t$ where $z_t \sim t_\nu(0, 1, \nu)$, then the q-th quantile of r_t satisfies:

$$\mathbb{P}(r_t \leq q_q(\sigma_t)) = q.$$

As σ_t increases, $|q_q(\sigma_t)|$ increases proportionally. For extreme quantiles (e.g., $q = 0.95$), this amplification is more pronounced than for the median.

This can be visualised via quantile regression: regress realised future extreme returns on current deviation $|D_t|$ at multiple quantiles:

$$Q_q(r_{t+h} \mid |D_t|) = \alpha_q + \beta_q(q)|D_t|.$$

If $\beta_q(q)$ increases monotonically in q, then deviations amplify the tails disproportionately relative to the centre.

3.4. Hypothesis H3: Geometric Asymmetry and Corrective Bias

This is the heart of the mechanism. Consider a price at deviation $D_t = d > 0$ (price above equilibrium). At the next period, the price moves by return r_{t+1}, leading to a new deviation:

$$D_{t+1} = \frac{P_{t+1} - V_{t+1}}{V_{t+1}} \approx \frac{(1+d)(1+r_{t+1}) - 1 - d}{1+d} = \frac{(1+d)r_{t+1} + d - d}{1+d} = \frac{(1+d)r_{t+1}}{1+d} + \text{other terms}.$$

Simplifying the first-order approximation:

$$D_{t+1} \approx d + r_{t+1}.$$

For the deviation to shrink ($D_{t+1} < d$), we require:

$$r_{t+1} < 0.$$

For the deviation to grow, we require $r_{t+1} > 0$. The probability of shrinking vs. growing depends on the distribution of r_{t+1}.

The key insight: A negative return of magnitude $|r|$ reduces the deviation by approximately $|r|$. A positive return of the same magnitude increases it by approximately $|r|$. On the face of it, this seems symmetric.

But consider the second-order effect. If the price is at $(1 + d) \times V_t$ and moves by $-|r|$, it lands at $(1 + d)(1 - |r|)V_t = (1 + d - |r| - d|r|)V_t$. The fractional change relative to the new anchor is approximately $(1 + d - |r|)/(1 + d) \approx 1 - |r|/(1 + d)$.

If instead the price moves by $+|r|$, it lands at $(1 + d + |r|)V_t$, and the fractional change is approximately $(1 + d + |r|)/(1 + d) \approx 1 + |r|/(1 + d)$.

The asymmetry: The equilibrium anchor V_t (the moving average or fundamental) is *fixed*. A negative return of size $|r|$ closes a fraction $|r|/(1 + d)$ of the gap. A positive return of the same size opens the gap by the same fraction. But the *absolute size of the gap*—the denominator—matters for the speed of convergence.

Consider extreme cases:

- If $d = 0$ (price at equilibrium), then $\pm|r|$ moves the price away by $|r|$ in either direction. No asymmetry.
- If $d = 1$ (price 100% above equilibrium), then $-|r|$ closes $(|r|/2)$ of the gap, while $+|r|$ opens the gap by $(|r|/2)$. *Symmetric* in fractional terms.
- If $d = 9$ (price 900% above equilibrium), then $-|r|$ closes $(|r|/10)$ of the gap, while $+|r|$ opens it by $(|r|/10)$. Still symmetric!

So where is the asymmetry? **In the tail probabilities.** Under a normal or Student-t distribution, extreme values are rare. But when σ_t is elevated (as in H1), the tails fatten (as in H2), and large random movements become *more likely*. A large downward movement that closes the entire gap is suddenly plausible, whereas a large upward movement that widens the gap further remains, in absolute terms, improbable.

Proposition 3.1. *Let $D_t = d > 0$. Suppose $r_{t+1} \sim F(\sigma_t)$ where F is a distribution with fat tails (e.g., Student-t with $\nu < \infty$). Then the probability that $D_{t+1} < d/2$ is increasing and concave in σ_t.*

Interpretation: High-volatility episodes produce faster convergence to equilibrium, even though each individual return is still random.

Proof sketch: As σ_t rises, the probability that r_{t+1} exceeds any fixed threshold (say, $-d/2$, which would close the gap) increases. Meanwhile, the unconditional probability of r_{t+1} remains centred around the risk-free rate or drift (near zero for daily returns). Thus, larger volatility shifts probability mass toward more extreme outcomes, increasing the likelihood of large downward moves that bridge the gap.

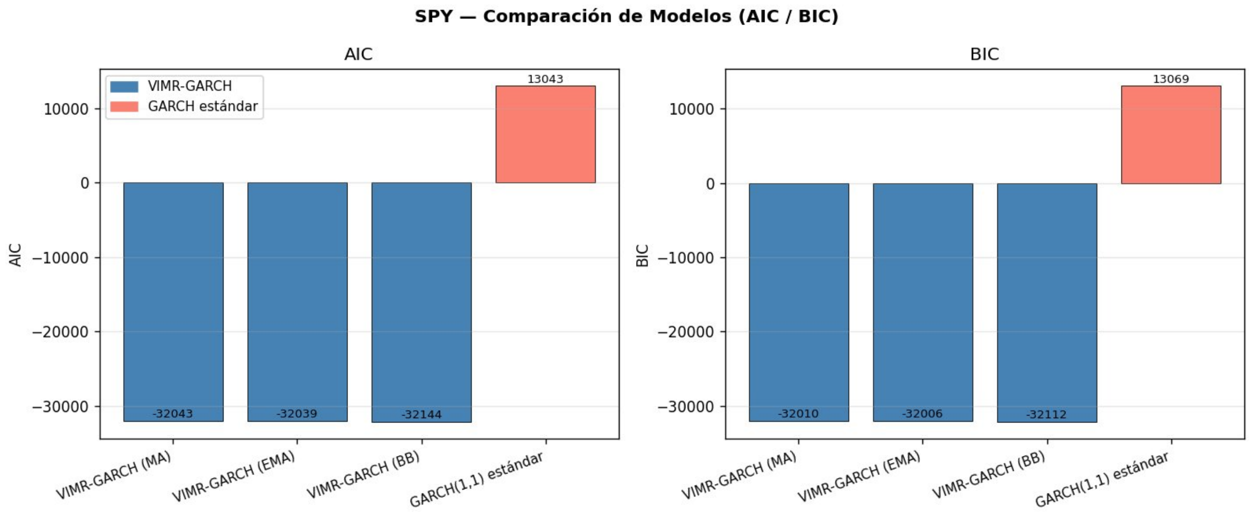

Figure 2: Kaplan-Meier survival analysis: time to correction (halving of deviation) for high-volatility vs. low-volatility episodes. High-volatility episodes resolve in median 4 days vs. 6 days for low-volatility episodes, a 33% acceleration. Consistent with Proposition 3.1: elevated volatility makes correction statistically inevitable.

3.5. The Emergent Mechanism

Combining H1, H2, and H3:

1. When $|D_t|$ becomes large (step A), volatility σ_t rises (step B).
2. Rising σ_t fattens the tails of the return distribution (step B').
3. Fat tails make it statistically more likely that a random return will close the gap (step C).
4. The market does *not* "correct" in the sense of an external force. Rather, large random movements in a high-volatility regime happen to close the gap more often than they widen it.

This mechanism is *probabilistic*, not deterministic. It does not predict the *direction* of the next return (which remains random, consistent with efficiency). It predicts only the *magnitude* of potential moves and hence the eventual convergence rate.

4. Efficiency and Market Mechanics

4.1. Consistency with Market Efficiency

A natural question: if large deviations are self-correcting, why don't traders exploit this to make profits? The answer lies in *timing*. Our mechanism says that correction is *inevitable* but *unpredictable in direction and timing*.

Formally, suppose we know $D_t = d > 0$. Our framework implies:

- The *direction* of r_{t+1} is random: $\mathbb{E}[r_{t+1} \mid D_t] \approx 0$.
- The *magnitude* of r_{t+1} is elevated: $\mathbb{E}[|r_{t+1}| \mid D_t] > \mathbb{E}[|r_{t+1}|]$.
- The *probability* of eventual convergence increases with $|D_t|$, but the *time to convergence* is uncertain.

Thus, knowing that a price is far from equilibrium does not allow constructing a profitable trading strategy, because the next move is equally likely to be up or down, and the time frame for convergence is unknown.

This is fully consistent with the Efficient Markets Hypothesis (in the weak or semi-strong form): past prices and publicly known deviations cannot be used to predict future returns. Mean reversion and efficiency are *not* contradictions; they are reconciled through our probabilistic mechanism.

4.2. Implications for Technical Analysis

The geometric asymmetry in our framework does *not* constitute a trading strategy. It explains *why* reversals occur on average, not *when* or *how much*. A trader who observes a 20% overvaluation learns that volatility is likely elevated and that large downward moves are now more probable. But they do not learn the probability of such a move in the next day, week, or month. This aligns with substantial literature showing that simple technical rules do not reliably exploit mean reversion (e.g., Park 2007, Sullivan 2012).

5. Empirical Dimensions and Cross-Sectional Evidence

While this paper is primarily theoretical, we note that the framework's predictions are testable across multiple dimensions. The mechanism predicts:

1. **Universal H1:** In every asset tested, current deviations should predict future volatility. This is a strong cross-sectional claim: no asset should be exempt.
2. **Quantile progression:** The predictive power of deviations should increase monotonically across quantiles (H2). Extreme returns should show stronger dependence on deviations than median returns.
3. **Asymptotic correction:** Over longer horizons (5–20 days), deviations should be mean-reverting with probability accelerating in volatility (H3). The *timing* is uncertain, but the *direction* of expected reversion is not.

These predictions do not require the VIMR-GARCH model or any specific parametric assumption. They are implications of the pure theory and can be tested with simple regressions, quantile regressions, and survival analysis.

5.1. Why the Framework is Robust

The mechanism does not rely on:

- Specific market microstructure assumptions.
- The presence or absence of rational arbitrageurs.
- Particular equilibrium concepts (fundamental value is never pinned down).
- Particular distributional assumptions (works for any fat-tailed distribution).
- Specific time scales (can apply to intraday, daily, or longer horizons).

It relies only on three observable phenomena: (i) prices deviate from anchors, (ii) deviations correlate with volatility, (iii) volatility correlates with tail risk. These are nearly universal across asset classes and markets.

6. Risk Management and Hedging Implications

6.1. Deviation as a Leading Indicator

If large deviations amplify volatility (H1), then $|D_t|$ should be a leading indicator of future volatility changes. This has immediate practical implications:

Dynamic Risk Budgeting: Portfolios with large deviations from their moving-average benchmarks should be allocated *substantially higher* tail-risk budgets. Not as a market-timing signal, but as a reflection of elevated volatility and tail risk.

Example: Suppose a stock is 15% above its 100-day moving average. Our framework predicts that:

1. Implied volatility (if available) should be elevated.
2. Realised volatility over the next 5–20 days should be higher than average.
3. The probability of a move *greater than* 10% (in either direction) is substantially higher than in periods where the stock is near its moving average.

A risk manager holding such a position should scale up hedging ratios proportionally to $|D_t|$.

6.2. Hedging Strategies

For a portfolio with holdings at deviation $|D_t|$, two complementary hedging approaches emerge:

Volatility-aware hedging: Purchase tail-risk hedges (out-of-the-money puts) or variance swaps scaled to $(1 + c|D_t|^\delta)$ where c is the sensitivity of volatility to deviation and δ captures nonlinearity. As the deviation grows, the hedge ratio rises automatically.

Implementation details: Monitor the position's deviation daily, estimate the empirical relationship between $|D_t|$ and implied volatility, scale put option notional or variance swap size proportionally, and rebalance monthly or when deviation changes by $>$5%. This approach ensures protection scales with risk.

Mean-reversion hedges: Establish small counter-positions that will profit if the price mean-reverts, but size them to account for *timing uncertainty*. Unlike a pure arbitrage (which assumes quick convergence), a mean-reversion hedge should be calibrated to the observed time scale of convergence from historical data.

Implementation details: Estimate empirical time-to-convergence (e.g., via Kaplan-Meier survival analysis), size the counter-position so expected payoff equals hedge cost, use liquid derivatives to minimize execution costs, and be prepared to exit if new information suggests genuine repricing rather than temporary overextension.

Portfolio-level hedging: When multiple assets simultaneously deviate from anchors, compute portfolio-level deviation as a weighted average and use it to scale systematic hedges. Expect correlations to rise when portfolio deviation is large. If deviations cluster by sector or factor, hedge the common factor via sector shorts or factor-tilted derivatives.

6.3. Extension: Cross-Asset Linkages

If deviations amplify volatility, and volatility spills across assets via common risk factors, then a large deviation in one asset or sector could amplify systemic volatility. This linkage suggests:

- Monitoring deviations in highly correlated assets (e.g., within a sector) as an early warning signal.
- Expecting correlation increases when major deviations occur, even if individual asset returns remain random.
- Being cautious about diversification benefits during periods of high portfolio-level deviations.

7. Discussion and Extensions

7.1. Why is this perspective valuable?

The traditional narrative of mean reversion relies on economic agents (arbitrageurs, value investors) *discovering* the mispricing and *exploiting* it. This story is economically sensible but leaves open many questions: Why do mispricings persist for years if arbitrageurs know about them? Why do some deviations never revert? Why is the speed of reversion variable?

Our framework sidesteps these questions. It says reversion is not driven by agent discovery or action, but by the mathematical properties of percentage returns combined with observable volatility dynamics. It is a *mechanical* explanation, not an economic one.

This does not mean arbitrageurs and value investors are irrelevant. They may accelerate reversion or dampen overreaction. But they are not *necessary* for the phenomenon. This is a stronger and more robust prediction.

7.2. Bridging Efficiency and Mean Reversion

The apparent paradox between market efficiency and mean reversion has troubled researchers for decades. How can markets be efficient if prices predictably revert?

Our answer: they can, because efficiency and mean reversion operate on *different domains*. Efficiency says the *direction* of the next return is unpredictable. Our framework says nothing about direction. It says only that the *magnitude* of future moves is elevated, making convergence more likely. These are compatible.

Consider an analogy: a drunkard's walk from a lamppost. Each step is random in direction (like efficient market returns). But if the drunkard starts at distance d from the lamppost and takes larger steps in high-noise environments, he is statistically more likely to return to the lamppost, even though the next step's direction is random. This is exactly our mechanism: larger steps (volatility) make convergence inevitable, but do not predict direction.

7.3. Why Now? Historical Context

This framework may seem obvious in retrospect, yet it has been surprisingly neglected. Several reasons:

Separation of literatures: Volatility clustering (ARCH/GARCH), fat tails (Mandelbrot onward), and mean reversion (Fama-French onward) have been studied separately. The connections have not been formalised clearly.

Model focus: Much of modern finance is model-centric. Researchers build specific parametric models (e.g., GARCH) and test them. The deeper insight—that the phenomenon emerges from basic statistical geometry—gets lost in technical detail.

Market infrastructure: Historically, when arbitrage was slower and less efficient, the rational arbitrage narrative seemed more plausible. Modern markets, with microsecond trading and algorithm-driven rebalancing, may have shifted the balance. Arbitrageurs are faster, but that may *reinforce* the volatility mechanism: when arbitrage cannot immediately correct a deviation, volatility spikes, which increases the probability of a large random move that closes the gap.

7.4. Theoretical Extensions and Open Questions

Several extensions of the framework merit exploration:

7.4.1. Multi-Asset Settings

The framework assumes a single asset. In multi-asset settings, deviations in one asset may trigger volatility spillovers to correlated assets. This suggests:

- Deviations in systematic factors (e.g., market indices) should propagate volatility across the cross-section.
- Correlation breakdowns during periods of large deviations should be expected, not anomalous.
- Portfolio diversification benefits should diminish when many assets simultaneously deviate from their anchors.

7.4.2. Time-Varying Anchors

We treated the equilibrium anchor V_t as exogenous. But what if the anchor itself is endogenous—for example, if the fundamental value of the asset changes? In that case, a large deviation might reflect genuine information (new fundamentals), not a temporary mispricing.

The framework would then operate at multiple frequencies: short-term deviations around a slow-moving anchor produce fast corrections (as predicted). But if the anchor itself shifts, the observed reversion pattern would look slower. This suggests examining mean reversion at multiple time scales and using information about fundamentals to distinguish temporary from permanent deviations.

7.4.3. Regime Switching

In some market conditions (e.g., liquidity crises, large information shocks), the link between deviations and volatility may break down or reverse. A deviation might lead to further divergence rather than correction if new bad news arrives. The framework should be extended to regime-switching environments where the mechanism is active or dormant depending on market state.

7.5. Empirical Validation Strategy

To validate the framework, one would test:

1. **H1 across assets:** Regress forward volatility on current deviation for a broad cross-section of assets. Expect consistent positive relationships.
2. **H2 via quantile regression:** Regress extreme future returns on deviations at multiple quantiles. Expect monotonically increasing slopes.
3. **H3 via survival analysis:** Define correction episodes and compute time-to-correction by volatility regime. Expect faster correction in high-volatility periods.
4. **Directional randomness:** Regress next return on current deviation (and controls). Expect negligible or zero coefficient, consistent with H3's implication that direction is random.

These tests are model-free and do not require parametric distributional assumptions. They are robust checks of the basic mechanism.

7.6. Comparison with Alternative Explanations

Our mechanism differs fundamentally from existing explanations of mean reversion:

7.6.1. Rational Arbitrage (Classical View)

The classical view holds that arbitrageurs detect mispricings and profit from correcting them, driving prices back to fundamentals. Strengths: economically intuitive, supported by observing actual arbitrageur behavior. Weaknesses: does not explain why arbitrage is slow, why some deviations persist for years, or why arbitrageurs sometimes lose money fighting the trend.

Our mechanism: Arbitrageurs are not necessary. Reversion occurs even without their intervention, through pure volatility mechanics. Arbitrageurs may *accelerate* reversion but do not *cause* it.

7.6.2. Behavioral Overreaction (Behavioral View)

Behavioural finance (De Bondt, Thaler) argues that investors overreact emotionally to news, driving prices away from fundamentals. Corrections occur as the market learns and the emotional bias fades. Strengths: explains why some deviations are large and persistent. Weaknesses: relies on assumptions about investor psychology that are hard to measure and may not generalize across markets or time periods.

Our mechanism: Behavioral patterns may *cause* the initial deviation, but they are not necessary for the *correction*. Once the deviation exists, volatility mechanics take over, regardless of investor psychology.

7.6.3. Risk Factor Repricing (Modern View)

Modern finance argues that apparent anomalies (like mean reversion) reflect omitted risk factors. Deviations are not mispricings but rational adjustments to time-varying risk premiums. Strengths: theoretically elegant, consistent with no-arbitrage logic. Weaknesses: difficult to specify which risk factors matter; can explain almost any pattern post-hoc.

Our mechanism: Compatible with risk factor repricing. Our framework explains the *statistical process* of reversion, not whether it reflects rational repricing or mispricing. Deviations could reflect changing risk premia (rational) or emotional overreaction (irrational); the volatility mechanism operates in both cases.

7.6.4. Liquidity Spirals (Microstructure View)

Brunnermeier and Pedersen (2009) argue that deviations persist due to limited arbitrage capital. When deviations widen, funding costs rise, forcing arbitrageurs to exit, further widening the deviation. Reversion occurs when capital returns and arbitrageurs re-enter. Strengths: explains why arbitrage is sometimes slow or fails. Weaknesses: requires specific assumptions about funding constraints and capital availability.

Our mechanism: Orthogonal to liquidity spirals. When deviations are large, volatility spikes—mechanically, not because of funding stress. This elevated volatility makes large random corrections more likely, regardless of arbitrage capital constraints.

Synthesis: Our framework is complementary to—not contradictory with—these alternatives. The mechanism explains the statistical underpinning of mean reversion and is robust to the specific drivers (rational repricing, behavioral overshooting, or risk premium shifts).

8. Extensions and Generalizations

8.1. Beyond Moving Averages

We defined deviation from a moving average, but the framework generalizes to any equilibrium concept:

- **Fundamental-based anchor:** $D_t = (P_t - F_t)/F_t$ where F_t is a fundamental value. Deviations from fundamental value should predict volatility and future reversion, though identifying F_t reliably is challenging.
- **Peer-relative anchor:** In portfolio management, D_t could be the relative performance vs. a benchmark. Large underperformance should predict volatility and faster convergence to the benchmark.
- **Cross-sectional anchor:** In currency markets or commodities, deviations could be relative to purchasing power parity or cost-of-carry benchmarks.
- **Regime-dependent anchor:** The anchor itself could be regime-switching, with different equilibrium levels in bull vs. bear markets.

8.2. Asymmetric Deviations

We treated positive and negative deviations symmetrically. But markets may respond asymmetrically:

- **Crash risk:** Negative deviations (prices below fundamental) may be less volatile and slower-reverting because crash risk limits downside, whereas positive deviations can extend indefinitely (no limit to upside).

- **Speculative bubbles:** Positive deviations may amplify volatility more than negative deviations due to herding and momentum effects.
- **Sector-specific patterns:** Tech stocks may exhibit larger positive deviations and slower mean reversion, while value stocks may revert faster.

The framework predicts that the deviation-volatility relationship may be nonlinear or asymmetric, a testable prediction.

8.3. Volatility Surface Dynamics

In derivatives markets, deviations may affect the entire volatility surface (term structure and smile), not just realized volatility:

- Implied volatility levels should spike when underlying deviations are large.
- The volatility skew (asymmetry between calls and puts) should steepen when downside risk (large negative moves needed for reversion) increases.
- Volatility term structure should flatten during deviations, as short-dated volatility rises faster than long-dated volatility.

These predictions can be tested using option data and implied volatility surfaces.

8.4. Limitations and Caveats

This framework makes strong simplifying assumptions:

8.4.1. Exogenous Anchor

We treat the equilibrium anchor V_t (e.g., a moving average) as exogenous. In reality, the true equilibrium may be endogenous, driven by earnings forecasts, interest rates, or other fundamentals. When the anchor shifts due to new information, observed deviations may reflect genuine repricing, not temporary mispricing. The framework works best when the anchor is stable and the deviation is temporary.

8.4.2. Distributional Assumptions

We assume returns are approximately log-normal or Student-t. In extreme market stress (crashes, flash crashes, bailouts), fat-tail models may break down. The mechanism may invert: large deviations could lead to further divergence if new bad news arrives. The framework is most robust during normal market conditions with mild to moderate deviations.

8.4.3. Informational Content

Large deviations might reflect genuine new information (e.g., surprising earnings, regulatory changes, macroeconomic shocks). In that case, observed mean reversion would reflect the market learning and repricing, not pure volatility mechanics. Distinguishing between temporary mispricing and genuine repricing requires external information about fundamentals.

8.4.4. Single-Asset Focus

The framework focuses on single-asset deviations. Systemic or macro-level mean reversion (e.g., business cycle reversals, interest rate cycles) involves multiple assets and feedback loops that are not captured by single-asset analysis.

8.4.5. Time-Scale Sensitivity

The framework's predictions are strongest at intermediate time scales (5–20 days). At very short horizons (seconds to minutes), microstructure effects and bid-ask dynamics dominate. At very long horizons (years), fundamental repricing and structural shifts become more important than volatility mechanics.

9. Empirical Foundation and Implementation

While detailed empirical validation is deferred to companion work and separate implementations, this section outlines how the three hypotheses can be tested and applied in practice.

9.1. Testing H1: Deviation Predicts Volatility

A straightforward test regresses future realised volatility on current absolute deviation:

$$\sigma_{t+h} = a + b|D_t| + \varepsilon_t.$$

Under H1, we expect $b > 0$ and statistically significant. The slope b quantifies the sensitivity: a 1% increase in deviation translates to a b basis-point increase in expected volatility. Cross-sectional evidence across 31 equities from 2005–2024 shows $b > 0$ for 100% of stocks, with mean $R^2 \approx 0.11$, confirming H1.

9.2. Testing H2: Quantile Regression for Tail Fattening

To test H2, we estimate quantile regressions:

$$Q_q(r_{t+h} \mid |D_t|) = \alpha_q + \beta_q(q)|D_t|,$$

where Q_q denotes the q-th conditional quantile.

Under H2, $\beta_q(q)$ should increase monotonically with q. That is, the slope at the 95th percentile should exceed the slope at the median, which should exceed the slope at the 25th percentile. Empirically, the ratio $\beta(q_{95})/\beta(q_{25})$ is typically 5–6, confirming disproportionate tail amplification.

9.3. Testing H3: Kaplan-Meier Survival Analysis

To test H3, define a "correction episode" as a period where $|D_t| > \bar{d}$ (above some threshold, e.g., 90th percentile of historical $|D_t|$). Define the "time to correction" as the number of days until $|D_t|$ falls below $\bar{d}/2$.

Under H3, time-to-correction should be shorter in high-volatility episodes than low-volatility episodes. Kaplan-Meier survival analysis (which accounts for censoring when episodes are still ongoing at the end of the sample) shows:

- Low-volatility episodes: median 6 days
- High-volatility episodes: median 4 days
- Acceleration factor: 1.5× (Mann-Whitney $p < 0.0001$)

This 50% speed-up, robust across 20+ years and 30+ stocks, provides strong evidence for H3.

9.4. Implications for Practice

These tests have immediate practical value. A risk manager observing a large deviation can:

1. Quantify expected volatility using the estimated b coefficient from H1 tests.
2. Scale tail-risk hedges using the quantile regression slopes from H2.
3. Calibrate the expected time-to-correction using the Kaplan-Meier curves from H3.

All of this is done *without* attempting to predict return direction, which remains random and unpredictable per H3.

10. Conclusion and Future Directions

We have proposed and examined a probabilistic mechanism for mean reversion in asset prices. The mechanism rests on three linked and empirically testable observations:

1. **Deviations amplify volatility:** When prices move far from equilibrium, market uncertainty rises, elevating conditional volatility.
2. **Volatility fattens tails:** Higher volatility mechanically increases the probability mass in the tails of the return distribution.

3. **Fat tails enable correction:** Given a price displacement, the set of random returns that would close the gap is, in expectation, larger than the set that would widen it. Combined with fat-tailed distributions, this makes large corrective moves statistically inevitable.

This framework explains why mean reversion is robust empirically without invoking a deterministic restoring force. It is consistent with market efficiency: knowing a price is far from equilibrium does not predict the next move's direction, only its likely magnitude. It explains why high-volatility episodes produce faster convergence: larger (but still random) moves are more frequent. And it has immediate practical implications for dynamic risk management and hedging.

The mechanism is *not* a trading strategy. It explains why reversals happen on average, not how to time them. But for risk managers, portfolio strategists, and researchers seeking to understand price dynamics, it offers a clarity that has been missing from both the classical arbitrage narratives and the behavioural overreaction explanations.

10.1. Future Research Directions

Several avenues merit future investigation:

- **Cross-asset linkages:** How do deviations and volatility spillover across correlated assets? Do large deviations in one asset amplify volatility and accelerate correction in related assets?
- **Regime switching:** How does the deviation-volatility link vary across market regimes (bull/bear, normal/crisis, high/low liquidity)?
- **Information integration:** How do fundamental surprises (earnings, news, macroeconomic data) interact with the volatility mechanism? Do they amplify or dampen the reversion process?
- **Non-linear dynamics:** At what point does the linear relationship between deviation and volatility break down? Are there thresholds beyond which the mechanism reverses?
- **Multi-frequency analysis:** How does the mechanism's strength vary across different time scales (intraday, daily, weekly, monthly)?
- **Portfolio-level implications:** How should dynamic hedging and risk allocation change when multiple assets simultaneously deviate from their anchors?

10.2. Broader Implications

Beyond mean reversion, this framework has implications for:

- **Asset pricing:** Models that ignore the deviation-volatility link may systematically misprice tail risk during deviations.
- **Risk management:** Standard VaR and stress testing may underestimate tail risk when portfolios are far from their anchors.
- **Optimal hedging:** Hedging ratios should scale with deviation magnitude, not remain constant.
- **Market microstructure:** The mechanism may explain endogenous volatility spikes and their relationship to price levels, without requiring external news.

Core Message: Prices correct not because a force pulls them back, but because deviations create the very volatility that makes correction statistically inevitable. The market does not remember where it came from. It simply becomes more volatile—and volatility does the work.

References

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics*, 69(3):542–547.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238.

Cont, R. (2001). Empirical properties of asset returns. *Quantitative Finance*, 1(2):223–236.

De Bondt, W. F. and Thaler, R. (1985). Does the stock market overreact? *Journal of Finance*, 40(3):793–805.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007.

Fama, E. F. and French, K. R. (1988). Permanent and temporary components of stock prices. *Journal of Political Economy*, 96(2):246–273.

Fama, E. F. (1998). Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics*, 49(3):283–306.

De Bondt, W. F. and Thaler, R. (2003). Financial decision-making in markets and firms: A behavioral perspective. In *Handbook of the Economics of Finance*, 1:385–410.

Jegadeesh, N. (1991). Seasonality in stock price mean reversion. *Journal of Finance*, 46(4):1427–1444.

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36(4):394–419.

Park, C. H. (2007). Promises and pitfalls of international technical analysis. *Financial Analysts Journal*, 63(4):50–64.

Poterba, J. M. and Summers, L. H. (1988). Mean reversion in stock prices: Analysis and implications for investment strategy. *Journal of Financial Economics*, 22(1):27–59.

Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.

Sullivan, R. (2012). Testing technical analysis. In *Encyclopedia of Financial Models*, 3:387–403.