

Volatility Absorption: The Diminishing Marginal Impact of Market Fear*

Yi-Hao Lai[†]

April 2026

Abstract

We document a novel empirical regularity: the marginal impact of fear shocks on realized asset returns diminishes as the ambient fear level rises—a phenomenon we term *volatility absorption*. Using daily data on four assets (SPY, GLD, TLT, 0050.TW) and VIX from 2006 to 2026, we measure absorption primarily through the *shock amplification ratio* (SAR)—the ratio of shock-day to normal-day absolute returns within each VIX regime—which removes VIX from the denominator of the dependent variable. We are explicit about the limit of that step: the SAR is still a ratio to a regime baseline that itself rises with VIX, so it measures multiplicative amplification relative to that baseline and does not by itself identify a declining marginal impact. The SAR declines from $3.15\times$ in calm markets ($VIX < 15$) to $2.33\times$ in high-fear regimes ($VIX\ 25\text{--}30$), with a slight reversal to $2.45\times$ in crisis regimes ($VIX \geq 30$); the calm-to-high decline of 0.82 carries a paired moving-block

*All errors are my own. **Replication.** Every number reported in this paper that is bound to an artifact is regenerated by a single entry point, `reproduce.py`, which reads the pinned 2026-04-19 price snapshot shipped with the package (`data/spy_gld_tlt_qqq_eem_vix_2005-2026.csv`) and the result files under `results/`, and exits non-zero on any mismatch; the current run matches 135 of 135 checks. The package will be deposited in full—snapshot, code, result files and the check log—at a public archive with a citable DOI on acceptance, and is available to referees on request during review. Two classes of number are deliberately *not* covered by that gate and are identified where they appear: the descriptive moments of 0050.TW, which the pinned snapshot does not contain, and the legacy shock-type, variance-risk-premium and hedging tables of earlier drafts, which are deferred extensions rather than results of this paper (Section 8).

[†]Department of Finance, Da-Yeh University. Corresponding author. Email: yhlai@mail.dyu.edu.tw

bootstrap p -value of 0.003. A contemporaneous GARCH null re-examination re-tires an earlier lagged-proxy simulation and shows that pooled null comparisons for this statistic are inconclusive—no single calibration of a GARCH null matches both the level and the daily increments of VIX—while a relative-threshold variant attributes roughly 58% of the pooled headline decline to fixed-threshold shock selection. The absorption claim therefore rests on a conditioning-and-sign decomposition: sorting days by the *pre-shock* ambient VIX level and restricting to genuine fear shocks ($\Delta\text{VIX} > +2$), the calm-to-high SAR decline is +1.05 with a 20-day paired moving-block bootstrap interval of $[0.33, 1.76]$ that excludes zero. Cross-asset regressions under the 2026-04-19 pinned snapshot show negative absorption coefficients for SPY, GLD, and TLT, with SPY only marginally significant and GLD/TLT significant at the 5% level. Event-level evidence from nonfarm payroll releases is directionally consistent, though we note the limitation of using a binary NFP dummy rather than surprise-magnitude controls. The paper is strongest as a reduced-form characterization of absorption and its robustness to alternative normalizers, while richer shock-type and portfolio-design extensions remain deferred pending fully reproducible supporting experiments.

Keywords: Volatility absorption, VIX, variance risk premium, fear shocks, endogenous risk, exogenous risk, volatility targeting, hedging

JEL Classification: G11, G12, G14, G17

Contents

1	Introduction	5
2	Literature Review	9
2.1	Volatility Clustering and GARCH Models	9
2.2	Variance Risk Premium	9
2.3	Leverage Effect and Asymmetric Volatility	10
2.4	Investor Attention, Habituation, and Behavioral Responses	12
2.5	Volatility Targeting	13
3	Methodology	13
3.1	VIX Regime Classification	13
3.2	Shock Amplification Ratio (Primary Measure)	14
3.3	Normalized Shock Impact (Supplementary Measure)	15
3.4	Absorption Regression	15
3.5	Shock Type Classification	16
3.6	Event-Level Analysis: Nonfarm Payrolls	17
3.7	Cross-Asset Tests	18
3.8	Variance Risk Premium Analysis	18
4	Data	18
4.1	Data Sources and Sample Period	18
4.2	Descriptive Statistics	19
4.3	VIX Regime Distribution	21
5	Results	22
5.1	The Absorption Effect	22
5.2	Cross-Asset Evidence	24
5.3	Event-Level Evidence: Nonfarm Payrolls	26
5.4	Shock-Type Extension (Deferred)	30
5.5	Variance Risk Premium Dynamics (Deferred)	30

6	Economic Implications	30
6.1	Hedging Cost-Benefit Analysis (Deferred)	31
6.2	Portfolio Rebalancing	31
6.3	Volatility Targeting Strategy Design	31
6.4	Crisis Response Timing	32
7	Robustness	33
7.1	Alternative Shock Thresholds	33
7.2	Sub-Period Analysis	34
7.3	Alternative Normalization	35
7.4	Controlling for Return Autocorrelation	36
7.5	Is the SAR Decline a Mechanical GARCH Artifact? A Contemporaneous Null Re-Examination	37
8	Conclusion	40
8.1	Limitations	42
8.2	Directions for Future Research	43
A	Variable Definitions	49
B	Additional Cross-Asset Results	49

1 Introduction

Financial markets exhibit a puzzling pattern during sustained crises: after an initial spike in volatility, subsequent shocks of comparable magnitude produce progressively smaller market reactions. During the 2008 Global Financial Crisis, for example, a two-point VIX increase when VIX was already at 60 moved SPY far less, in volatility-adjusted terms, than the same two-point increase when VIX was at 15. Market participants, policymakers, and risk managers have long observed this “numbness” informally, yet the phenomenon has received surprisingly little systematic empirical attention.

We formalize and quantify this observation under the label *volatility absorption*: the diminishing marginal impact of fear shocks on realized asset price movements as the ambient level of market fear rises. Our primary measure is the *shock amplification ratio* (SAR)—the ratio of mean absolute returns on shock days to non-shock days *within* each VIX regime—which removes VIX from the denominator of the dependent variable and so avoids the specific artifact that afflicts VIX-normalized measures. It does not, however, make the statistic free of denominators altogether, and Section 5.1 sets out what remains: the SAR’s own denominator is the regime baseline, which rises with VIX, so a constant absolute shock increment already implies a declining SAR. When VIX is below 15 (“calm”), a large VIX shock ($|\Delta\text{VIX}| > 2$) produces an absolute SPY return that is 3.15 times the normal-day average. This ratio declines to $2.77\times$ (VIX 15–20), $2.38\times$ (VIX 20–25), and $2.33\times$ (VIX 25–30), with a slight uptick to $2.45\times$ in crisis markets ($\text{VIX} \geq 30$). The overall calm-to-high decline of 0.82 is economically large and statistically significant under a paired moving-block bootstrap ($p = 0.003$; Table 3). Two qualifications discipline how much of this pooled pattern we attribute to absorption. First, the pooled statistic mixes fear shocks ($\Delta\text{VIX} > +2$) with relief rallies ($\Delta\text{VIX} < -2$), and part of the decline reflects fixed-threshold shock selection rather than economics: switching to a relative threshold cuts the headline decline by roughly 58% (Section 7.5). Second, and constructively, the mechanism-relevant version of the question—does a genuine fear shock move prices less when *pre-existing* ambient fear is already high?—survives: conditioning regimes on the pre-shock VIX level and restricting to upward VIX shocks, the calm-to-high SAR decline

is +1.05 with a paired moving-block bootstrap interval that excludes zero (Section 7.5).

As a supplementary continuous-variable summary, we also compute the *normalized shock impact* ($\text{NSI} = |r_t|/V_t$) and regress it on VIX level. The slope is -0.000267 ($t = -1.77$, $p = 0.076$) under the 2026-04-19 pinned snapshot ($N = 768$), directionally consistent with the absorption hypothesis.¹ We emphasize that this “paralysis” coefficient is consistent with—but does not substitute for—the SAR evidence, because the NSI construction introduces a mechanical negative correlation through the VIX denominator. Identification therefore rests on the SAR, which does not divide by VIX; the shock-type decomposition that earlier drafts used as a second, denominator-invariant argument is withdrawn in this revision (Section 5.4) and is no longer offered in support of it.

An earlier draft emphasized a decomposition of absorption by shock type. That classification remains economically interesting, but the archived counts and bootstrap test statistics are not yet reproducible under the current pinned-snapshot workflow, so we no longer treat it as headline evidence in this version of the paper. The conceptual distinction between endogenous risks already reflected in VIX and exogenous risks that arrive outside the existing fear state remains a useful hypothesis for future work, especially given Danielsson et al. (2018)’s emphasis on endogenous risk dynamics within financial crises.

Event-level evidence provides directionally supportive evidence. We examine nonfarm payroll (NFP) release days—a scheduled macroeconomic event whose impact should, in theory, depend on the surprise component rather than the ambient fear level. Overall, NFP days produce absolute returns 1.16 times the all-non-NFP average (unadjusted $p = 0.039$; $1.19\times$ versus Fridays only, unadjusted $p = 0.036$; Welch’s two-sample t -tests, $N_{\text{NFP}} = 194$, sample 2010–2026), but we do not claim significance for this: Holm correction across the two comparisons adjusts both to $p = 0.072$, and the equal-variance

¹Under the paper’s original data snapshot (2025 yfinance pull), the corresponding estimate was -0.00028 ($t = -3.42$). The sign is preserved but the t -statistic weakens substantially due to yfinance retroactive adjustment of VIX historical data, which changes the shock-day classification and regression N from 893 to 768. The RV-normalized alternative ($\hat{\beta}^{\text{RV}} = -0.01249$, $t = -8.2$) is unaffected by this revision and provides the stronger quantitative evidence. See Section 7 for full snapshot-sensitivity disclosure.

variant of the vs.-all test gives $p = 0.051$.² In calm markets ($VIX < 15$) the ratio is $1.31\times$ (unadjusted $p = 0.026$, $n = 63$); in the medium-VIX regime (15–20) it is $1.23\times$ ($p = 0.029$). When VIX exceeds 25 the measured effect disappears, the ratio falling to $0.94\times$ ($p = 0.707$), though with $n = 28$ that cell is too imprecise to establish a reversal. The ordered decline from low to high VIX is consistent with the absorption hypothesis, but we do not claim it as established, for two reasons. The regime p -values are unadjusted and none survives Holm correction (smallest adjusted $p = 0.104$); and tested directly, the calm-minus-crisis ratio difference is $+0.37$ with a 95% bootstrap interval of $[-0.10, 0.79]$ that includes zero. The NFP evidence is descriptive support, not the basis of the paper’s inference.

We extend our analysis to multiple assets and find negative normalized regression slopes in three of the four assets we examine, significant at the 5% level for two of them. SPY (slope -0.000273), GLD (-0.000434), and TLT (-0.000437) all exhibit negative normalized regression slopes under the 2026-04-19 pinned snapshot, confirming diminishing marginal fear impact (see Table 9 and Section 7 for snapshot-sensitivity disclosure). The sole exception is the Taiwan equity ETF 0050.TW (slope $+0.000092$), which we interpret as evidence that VIX is a US-specific fear gauge with limited direct relevance for non-US equity markets. Taiwanese equities respond to VIX indirectly through US lead-lag effects, and a US-specific VIX shock may not capture the incremental fear relevant to the Taiwan market.

Our findings connect to several strands of the literature. First, the volatility clustering literature (Engle, 1982; Bollerslev, 1986) documents that high-volatility periods persist, and the news impact curve framework of Engle and Ng (1993) parameterizes how shock magnitude varies with prior shocks, but neither addresses the diminishing *marginal*

²The overall NFP volatility ratios and all regime-specific ratios reported in this section are estimated from 194 NFP release dates over 4,084 SPY trading days (194 NFP plus 3,890 non-NFP days, January 2010 to March 2026). Release dates are taken from the official Bureau of Labor Statistics “Employment Situation” release calendar, retrieved from the Federal Reserve Bank of St. Louis ALFRED real-time archive (release id 50); we do not use a first-Friday-of-month approximation, which misdates 33 of the 194 releases in this window and additionally assigns an event to October 2025, a month in which no Employment Situation report was published. Where a release falls on a market holiday—five Good Fridays in this sample, on which BLS published but US equity markets were closed—the event is assigned forward to the next trading day, since no reaction is observable on the release date itself. Data source: yfinance (SPY, ^VIX); reproducible values are bound in `paper/volatility-absorption/reproduce.py`.

impact within sustained high-volatility episodes. The threshold GARCH model of [Zakoian \(1994\)](#) allows for state-dependent shock responses, making it the closest parametric analogue to our nonparametric findings. Second, the variance risk premium (VRP) literature ([Bollerslev et al., 2009](#); [Carr and Wu, 2009](#); [Todorov, 2010](#)) shows that implied volatility systematically exceeds realized volatility on average. This paper takes no empirical stand on how the VRP behaves across VIX regimes: our earlier regime-level VRP estimates are not reproducible at the standard applied to the retained evidence, and a surviving legacy artifact points to VRP sign flips in crisis regimes when taken literally, so the VRP analysis is deferred in this revision (Section 5.5). Third, the literature on the leverage effect ([Black, 1976](#); [Christie, 1982](#)) documents asymmetric volatility responses to returns but does not examine how the magnitude of this response varies with the ambient fear level. Fourth, the investor attention literature—including work on the VIX as an “investor fear gauge” ([Whaley, 2000](#)), the FEARS index ([Da et al., 2015](#)), and attention-based models ([Andrei and Hasler, 2015](#))—suggests that investors’ information-processing capacity becomes saturated during sustained crises, providing a behavioral foundation for our empirical findings.

The economic implications of volatility absorption are substantial, but this draft now states them more cautiously. If shocks have diminishing marginal impact at high VIX, then the incremental value of adding protection or reacting mechanically to every additional fear spike should be lower in already-stressed regimes. This logic is consistent with volatility-targeting frameworks that scale exposure inversely with VIX, though the paper does not rely on unreproduced portfolio-backtest numbers to make that point.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes our methodology. Section 4 presents the data. Section 5 reports the main results. Section 6 discusses economic implications. Section 7 presents robustness checks. Section 8 concludes.

2 Literature Review

Our paper relates to four main streams of the finance literature: (i) volatility clustering and GARCH modeling, (ii) the variance risk premium, (iii) the leverage effect and asymmetric volatility, and (iv) investor attention and fear habituation.

2.1 Volatility Clustering and GARCH Models

The observation that “large changes tend to be followed by large changes” was first formalized by [Mandelbrot \(1963\)](#) and later embedded in the autoregressive conditional heteroskedasticity (ARCH) framework of [Engle \(1982\)](#) and its generalized extension (GARCH) by [Bollerslev \(1986\)](#). These models capture the persistence of volatility but treat each shock’s impact as governed by fixed parameters (α, β, γ) . The threshold GARCH (TGARCH) model of [Zakoian \(1994\)](#) and its extensions allow the shock impact to differ across sign regimes (positive vs. negative returns), while [Engle and Ng \(1993\)](#) develop the *news impact curve* framework that parameterizes how the magnitude of a shock’s impact on future volatility varies with both the sign and size of the shock. Our work is related to but distinct from these frameworks: we do not model sign-dependent or size-dependent shock impacts within a GARCH structure, but rather document that the *effective* impact of a shock on *contemporaneous returns* depends on the ambient volatility regime. Regime-switching GARCH models ([Hamilton, 1989](#); [Haas et al., 2004](#)) implicitly allow for regime-dependent dynamics, but they treat each regime as having fixed parameters. Our contribution is to show that the marginal impact of shocks declines *within* the high-volatility regime itself, a nonlinearity that standard regime-switching models do not capture. See also [Muler and Yohai \(2008\)](#) for robust GARCH estimation under the type of heavy-tailed outliers that characterize our crisis-regime observations.

2.2 Variance Risk Premium

[Bollerslev et al. \(2009\)](#) document a significant variance risk premium (VRP)—the spread between risk-neutral and physical expected variance—and show that it predicts future eq-

uity returns. Carr and Wu (2009) provide a model-free approach to measuring VRP using options. Todorov (2010) decomposes the VRP into continuous and jump components, showing that jump risk commands a substantial premium, particularly during periods of market stress. Drechsler and Yaron (2011) develop a long-run risk model linking the VRP to macroeconomic uncertainty, while Bekaert et al. (2022) document substantial time variation in the VRP tied to risk appetite and macro uncertainty. More recently, Bekaert and Hoerova (2014) decompose the VIX² into expected variance and the variance risk premium, showing that the premium component drives most of the predictive power. We cite this literature as motivation for why conditioning on the VIX level is economically meaningful—the VIX embeds a time-varying premium component, not just a physical variance forecast—rather than as a set of findings we extend: our own regime-level VRP analysis is deferred in this revision (Section 5.5). Martin (2017) provides an alternative framework for linking expected returns to option-implied variance that may offer additional insights into the VRP dynamics we document.

2.3 Leverage Effect and Asymmetric Volatility

Black (1976) first noted that volatility tends to rise more after negative returns than after positive returns of equal magnitude. Christie (1982) attributed this to financial leverage, while Bekaert and Wu (2000) emphasized the role of time-varying risk premiums. The GJR-GARCH model of Glosten et al. (1993) and the EGARCH model of Nelson (1991) were developed to capture this asymmetry. Our contribution is orthogonal: we do not examine the direction-dependent response to returns but rather the *level-dependent* response to volatility shocks. Volatility absorption implies that the effective α (shock impact) in a GARCH framework is not constant but decreases with the level of conditional variance—a form of nonlinearity that standard GARCH models do not capture.

A closely related strand studies the contemporaneous relation between index returns and changes in option-implied volatility directly. Fleming et al. (1995) analyzed the CBOE market volatility index then in dissemination (computed from S&P 100 options, the series now labelled VXO) and documented that changes in implied volatility are

strongly and negatively associated with contemporaneous index returns. [Low \(2004\)](#) shows that this risk–return relation is neither linear nor symmetric, and is best described as a reclined S-curve: extreme price declines coincide with rapid upward spikes in implied volatility, while extreme price advances produce comparatively muted declines, an asymmetry that [Low \(2004\)](#) interprets as a further form of loss aversion. [Hibbert et al. \(2008\)](#) confirm a strong negative return–implied-volatility relation at both daily and intra-day frequencies, find neither the leverage nor the volatility-feedback hypothesis sufficient to explain it, and attribute the pattern to representativeness, affect, and extrapolation bias; in their evidence the magnitude of the relation and its asymmetry are most closely associated with *extreme changes in index returns*. [Fousekis \(2020\)](#) revisits the shape nonparametrically across four stock–implied-volatility index pairs and likewise finds a negative, contemporaneous, nonlinear association that is asymmetric in both the sign and the size of the return, while reading the estimated reverse S-shape as sitting uneasily with the fear-and-exuberance interpretation.

Our design differs from this strand in *what indexes the nonlinearity*. There, the nonlinearity is a function of the return shock’s own sign and magnitude: the relation bends as the move grows larger or changes sign. Here, the conditioning variable is instead the *ambient* level of implied volatility prevailing before the shock arrives. The SAR compares shock days against non-shock days drawn from the same fear environment, and the object of interest is how that within-regime amplification changes as the environment itself becomes more fearful. We flag at the outset that the headline SAR of Table 3 bins on the *same-day* VIX level and is therefore subject to post-shock sorting; it is the ambient specification of Section 7.5—regimes defined on the pre-shock level V_{t-1} while the sample is restricted to upward VIX shocks—that isolates the ambient axis, and that specification carries the paper’s claim. In that decomposition the sign asymmetry documented by [Low \(2004\)](#) and [Fousekis \(2020\)](#) is held fixed, and what remains varying is the ambient level; neither axis subsumes the other.

The ambient axis is not unexamined in the wider literature, and we do not claim it as unmeasured. [Just and Echaust \(2020\)](#) estimate a two-regime Markov switching

model for S&P 500 returns on changes in VIX and report a coefficient of -0.0897 in the low-volatility regime against -0.2405 in the high-volatility regime—roughly a three-fold *higher* sensitivity when ambient volatility is high, the opposite direction to the pattern we document. That contrast is instructive rather than contradictory, because the estimands differ: [Just and Echaust \(2020\)](#) measure the elasticity of the return level to a contemporaneous VIX change with regimes inferred from the return process itself, whereas the SAR measures the ratio of mean absolute returns on shock days to non-shock days within an exogenously binned pre-shock VIX level. Reconciling the two—whether a steeper regime-specific slope is compatible with a smaller shock-day amplification relative to an elevated baseline—is an open question that the present design does not settle, and we flag it as the sharpest external challenge to the interpretation offered here.

2.4 Investor Attention, Habituation, and Behavioral Responses

The behavioral basis for volatility absorption draws on several distinct literatures. First, the concept of diminishing sensitivity in prospect theory ([Kahneman and Tversky, 1979](#))—the concavity (convexity) of the value function for gains (losses)—provides a general psychological foundation for why successive units of bad news may have declining marginal impact, though we note that prospect theory concerns gains and losses relative to a reference point rather than shock magnitude relative to ambient volatility. [Barberis et al. \(2001\)](#) formalize a model of investor sentiment in which prior outcomes affect risk tolerance, generating momentum and mean-reversion patterns that are conceptually related to—but distinct from—the absorption phenomenon we document. Second, the investor attention literature provides a more direct channel. [Da et al. \(2015\)](#) construct the FEARS index from Google search volume for anxiety-related terms and show that investor sentiment directly affects asset prices. [Vlastakis and Markellos \(2012\)](#) demonstrate that information demand (proxied by search volume) covaries with stock market volatility, providing a JBF-published framework for understanding how attention interacts with fear. [Andrei and Hasler \(2015\)](#) model investor attention as endogenous to market conditions, showing that attention is higher during volatile periods. We interpret

our absorption findings as consistent with a specific prediction of attention-based models: when markets are already in a heightened state of fear (high VIX), the marginal informational surprise from an additional fear shock is lower because investors are already highly attentive and have already incorporated much of the crisis-relevant information into prices. This mechanism is distinct from the “habituation” to repeated stimuli documented in the psychology literature, and we avoid attributing it to prospect theory directly.

2.5 Volatility Targeting

Volatility targeting strategies, which scale risky-asset exposure inversely with predicted volatility, have been studied extensively. [Moreira and Muir \(2017\)](#) show that VT strategies improve the Sharpe ratio of equity portfolios, though the mechanism is debated. [Harvey et al. \(2018\)](#) provide a comprehensive analysis of VT performance and attribute the gains primarily to drawdown reduction rather than return enhancement. [Fleming et al. \(2001\)](#) demonstrate that volatility timing with daily rebalancing improves portfolio performance. Our contribution connects VT to the absorption framework: VT strategies with a c/VIX weight function naturally account for absorption because they reduce exposure precisely when the marginal impact of shocks is declining.

3 Methodology

3.1 VIX Regime Classification

We partition trading days into five VIX regimes to capture the granular structure of the absorption effect:

- **Calm** ($V_t < 15$): Low-fear environment, typically associated with steady markets.
- **Normal** ($15 \leq V_t < 20$): Baseline environment, the most common regime.
- **Elevated** ($20 \leq V_t < 25$): Moderate-fear environment, often surrounding earnings seasons or policy announcements.

- **High** ($25 \leq V_t < 30$): High-fear environment, typically during significant market stress.
- **Crisis** ($V_t \geq 30$): Extreme-fear environment, associated with major crises (GFC, COVID).

Let r_t denote the daily log return of an asset and V_t denote the VIX closing level on day t . We define a *fear shock* as a day on which $|\Delta V_t| > \tau$, where τ is a threshold (our baseline uses $\tau = 2$).

3.2 Shock Amplification Ratio (Primary Measure)

Our primary measure of absorption is the *shock amplification ratio* (SAR). For each VIX regime j , the SAR compares the mean absolute return on shock days to the mean absolute return on non-shock days *within the same regime*:

$$\text{SAR}_j = \frac{\overline{|r_t|}_{t \in S_j}}{\overline{|r_t|}_{t \in N_j}} \quad (1)$$

where S_j is the set of shock days in regime j and N_j is the set of non-shock days in regime j . The SAR compares shock-day to non-shock-day absolute returns *within* each VIX bin and does not divide by VIX, so it is free of the artifact that arises when the regressor appears in the denominator of the dependent variable. It is not, however, free of denominators. Writing $b_j = \overline{|r|}_{\text{normal},j}$ for the regime baseline, a shock that adds the same absolute increment δ to returns in every regime gives $\text{SAR}_j = (b_j + \delta)/b_j = 1 + \delta/b_j$, and since b_j rises with the VIX regime, SAR_j falls across regimes with no change in the shock’s marginal effect whatsoever. Volatility absorption predicts a decline in SAR across regimes, but so does that null, and the two are not separated by any estimate reported in this paper. We therefore treat a declining SAR as evidence of declining *multiplicative* amplification relative to a rising baseline, and we do not claim it establishes a declining marginal impact. The specification that would separate them—a conditional-response model $|r_t| = m(V_{t-1}) + g(\text{shock size}, V_{t-1}) + \varepsilon_t$ estimating g directly against the ambient level, with no normalization—is not estimated here and is identified in Section 8 as this

paper’s principal outstanding task. Perfect monotonicity is also not predicted; Section 5.1 reports where it fails.

3.3 Normalized Shock Impact (Supplementary Measure)

As a supplementary continuous-variable measure, we define the *normalized shock impact*:

$$\text{NSI}_t = \frac{|r_t|}{V_t} \quad (2)$$

The NSI provides a convenient scalar summary of shock intensity per unit of ambient fear, but it carries an important identification caveat.

Mechanical correlation caveat. Because VIX appears in the denominator of NSI and simultaneously defines the conditioning regimes, a purely mechanical negative relationship can arise: even if $|r_t|$ were statistically independent of V_t , dividing by a larger V_t would produce a smaller NSI, generating a spurious declining NSI–VIX pattern. The SAR (Equation 1), which does *not* divide by VIX, is what we offer against this concern: it independently confirms the declining pattern (Section 5.1). We are explicit that the two further arguments earlier drafts made here no longer stand. The endogenous/exogenous decomposition is withdrawn in this revision (Section 5.4) and cannot support any claim. The realized-volatility normalizer (Section 7) moves in the same direction, but it shares the denominator–regressor correlation structure that is at issue, so it is a cross-normalizer consistency check and not an independent refutation of the mechanical channel. The NSI regression provides a useful continuous summary statistic, but our identification rests on the SAR alone.

3.4 Absorption Regression

To formalize the relationship using the supplementary NSI measure, we estimate the cross-sectional regression:

$$\text{NSI}_t = \alpha + \beta \cdot V_t + \varepsilon_t, \quad t \in \{t : |\Delta V_t| > \tau\} \quad (3)$$

A negative β (the “absorption coefficient”) confirms that the marginal impact of fear shocks diminishes with the ambient fear level. We report Newey-West standard errors with 10 lags to account for heteroskedasticity and autocorrelation. One caveat on the lag choice: the regression sample consists of shock days only, so consecutive *observations* are not consecutive *calendar* days—except in crisis episodes, where shocks cluster and adjacent observations genuinely are adjacent days. Ten lags in observation time therefore span a variable and regime-dependent calendar window; we choose 10 as a conservative bracket for the shock-clustering episodes that dominate the autocorrelation, and note that the primary SAR inference (Table 3) does not depend on this choice, as it uses a moving-block bootstrap on the full daily series instead.

3.5 Shock Type Classification

We classify negative-return shock days ($r_t^{\text{SPY}} < 0$ and $|\Delta V_t| > \tau$) into three *mutually exclusive* types using a priority ordering based on the joint behavior of bonds and gold. Because some shock days satisfy multiple criteria simultaneously (e.g., $\text{SPY}\downarrow$, $\text{TLT}\uparrow$, $\text{GLD}\uparrow$ could qualify as both risk-off and geopolitical), we impose the following sequential classification rule to ensure each day is assigned to exactly one category:

1. **Geopolitical shocks** (highest priority): $r_t^{\text{SPY}} < 0$ and $r_t^{\text{GLD}} > 0.5\%$. Gold rallies meaningfully as a geopolitical safe haven while equities decline. The 0.5% threshold isolates days with economically significant gold moves rather than incidental co-movement. This category is checked first because geopolitical risk represents genuinely exogenous information not captured by bond-equity dynamics.
2. **Risk-off flights**: $r_t^{\text{SPY}} < 0$ and $r_t^{\text{TLT}} > 0$ (and not already classified as geopolitical). Classic flight-to-quality: money flows from equities to government bonds.
3. **Rate shocks** (residual): $r_t^{\text{SPY}} < 0$ and $r_t^{\text{TLT}} < 0$ (and not already classified above). Both equities and bonds fall, typically driven by interest rate surprises or inflation fears.

Days with $r_t^{\text{SPY}} < 0$ and $|\Delta V_t| > \tau$ that do not satisfy any of the above criteria (e.g., $r_t^{\text{TLT}} = 0$ and $r_t^{\text{GLD}} < 0.5\%$) are excluded from the shock-type analysis. The priority ordering ensures that the geopolitical category captures the most exogenous shocks first, before the remaining days are split between endogenous risk categories.

For each shock type k , we compute the type-specific absorption coefficient:

$$\text{Absorption}_k = \overline{\text{NSI}}_{\text{calm},k} - \overline{\text{NSI}}_{\text{high},k} \quad (4)$$

A positive value indicates that the normalized impact *declines* from calm to high VIX—i.e., the shock is absorbed. A negative or near-zero value indicates that the shock retains its full normalized impact regardless of the ambient fear level (no absorption).

Hypothesis 1. *Endogenous absorption. Risks that are endogenous to the information set priced by VIX—such as rate shocks and risk-off flights—are absorbed (positive absorption coefficient), because VIX already reflects these risks. Exogenous risks—such as geopolitical shocks—are not absorbed, because they represent genuinely new information orthogonal to the prevailing fear level.*

3.6 Event-Level Analysis: Nonfarm Payrolls

To provide event-level evidence, we follow the macro-announcement event-study tradition (Andersen et al., 2003; Balduzzi et al., 2001) and identify all US nonfarm payroll release dates over our sample period using the Bureau of Labor Statistics calendar. We note that the standard approach in this literature uses the *surprise component* (actual minus consensus forecast) as the key regressor; our analysis uses a binary NFP dummy due to data constraints, and we discuss the implications of this limitation in Section 5. For each NFP day, we compute:

$$\text{NFP ratio}_j = \frac{\overline{r_t^{\text{SPY}}|_{t \in \text{NFP} \cap j}}}{\overline{r_t^{\text{SPY}}|_{t \in \text{non-NFP} \cap j}}} \quad (5)$$

where j indexes VIX regimes. Under the absorption hypothesis, $\text{NFP ratio}_{\text{high}} < \text{NFP ratio}_{\text{calm}}$: the scheduled macro event produces less incremental volatility when VIX is already elevated.

3.7 Cross-Asset Tests

We replicate the absorption regression (Equation 3) for each asset in our sample: SPY, GLD, TLT, and 0050.TW. Because VIX is a US equity options-implied measure, we expect the absorption effect to be strongest for US equities and weaker for assets whose returns are less directly driven by US equity-market fear.

3.8 Variance Risk Premium Analysis

We compute the realized variance using the standard estimator:

$$RV_t^{(h)} = \sum_{i=1}^h r_{t-h+i}^2 \quad (6)$$

where $h = 22$ trading days (one month). The variance risk premium is:

$$VRP_t = VIX_t^2/252 - RV_t^{(22)}/22 \quad (7)$$

annualized to a daily scale. We examine how the VRP varies across VIX regimes.

4 Data

4.1 Data Sources and Sample Period

Our primary dataset consists of daily closing prices for four assets and one volatility index, all obtained from Yahoo Finance (yfinance API) and the Chicago Board Options Exchange (CBOE):

- **SPY**: SPDR S&P 500 ETF Trust, the most liquid US equity ETF, serving as our primary equity proxy.
- **GLD**: SPDR Gold Shares ETF, representing gold as an alternative safe haven asset.
- **TLT**: iShares 20+ Year Treasury Bond ETF, representing long-duration US government bonds.

- **0050.TW**: Yuanta/P-shares Taiwan Top 50 ETF, representing the Taiwanese equity market.
- **VIX**: CBOE Volatility Index, computed from S&P 500 index options, our primary measure of market fear.

The sample period spans January 2006 to early April 2026. Under the 2026-04-19 pinned snapshot used for all headline statistics, the SPY–VIX panel contains 5,092 trading days with joint availability of the SPY return, the VIX level, and the VIX change (2006-01-05 to 2026-04-02); 0050.TW has fewer observations due to its trading calendar. This period encompasses multiple market regimes: the pre-crisis expansion (2006–2007), the Global Financial Crisis (2008–2009), the European debt crisis (2011–2012), the low-volatility regime (2013–2017), the COVID-19 crash (February–March 2020), the post-COVID recovery and inflation episode (2021–2023), and the monetary tightening cycle (2022–2025).

4.2 Descriptive Statistics

Table 1 reports summary statistics for daily returns and VIX levels, recomputed for this revision directly from the pinned snapshot rather than carried over from the earlier draft (see the table note). SPY has a mean daily return of approximately 0.04% with a standard deviation of 1.22%, consistent with the well-documented equity premium. GLD and TLT exhibit comparable volatility, with TLT the least volatile of the three. VIX averages approximately 19.5 over the full sample, with a minimum of 9.14 (November 2017) and a maximum of 82.69 (March 2020).

Table 1: Descriptive Statistics: Daily Returns and VIX (2006–2026)

	Mean	Std Dev	Skewness	Kurtosis	Min	Max	N
<i>Panel A: Daily Returns (%)</i>							
SPY	0.040	1.219	−0.30	14.3	−11.59	13.56	5,094
GLD	0.042	1.159	−0.46	6.8	−10.84	10.70	5,094
TLT	0.012	0.935	0.01	3.3	−6.90	7.25	5,094
<i>Panel B: VIX Level</i>							
VIX	19.48	8.70	2.51	9.4	9.14	82.69	5,094

Notes: Returns are computed as log returns $r_t = \ln(P_t/P_{t-1}) \times 100$ on adjusted closes. Kurtosis is reported in excess form. VIX is reported in level form. Sample period: 2006-01-01 to 2026-04-02 under the 2026-04-19 pinned snapshot. Data source: Yahoo Finance (yfinance API) and CBOE.

Provenance of this table, revised. Every cell above is recomputed from the paper’s own pinned CSV (data/spy_gld_tlt_qqq_eem_vix_2005-2026.csv)

by scripts/rebuild_desc_stats.py and bound to results/desc_stats.json. Earlier drafts of this table reported $N = 5,050$ (4,850 for 0050.TW) together with moments that we can no longer reproduce from any surviving artifact: for SPY the earlier draft gave std 1.194, skewness −0.42 and max 9.38 against the 1.219, −0.30 and 13.56 obtained here. We report the recomputed values and withdraw the earlier ones rather than reconciling them, because the original computation left no record and we therefore cannot establish what it measured. Two further discrepancies are disclosed rather than resolved. First, the $N = 5,094$ here exceeds the $N = 5,092$ reported for the absorption regressions (Table 2) by two days. The two figures come from different objects: this table is computed from the pinned CSV, whose first row in the window is 2006-01-03, whereas the regression artifacts are computed from the 2026-04-19 yfinance snapshot and run from 2006-01-05. We have not established which two days account for the gap, and requiring a same-day VIX change does not explain it—on the pinned CSV that restriction removes no observations. Second, the first-day difference means the two samples are not nested, so the descriptive moments here are

4.3 VIX Regime Distribution

Table 2 reports the distribution of trading days across the five VIX regimes under the pinned snapshot. Calm markets ($VIX < 15$) account for approximately 34% of trading days, normal markets (15–20) for 31%, elevated markets (20–25) for 17%, high-fear markets (25–30) for 9%, and crisis markets ($VIX \geq 30$) for 9%. The shock rate ($|\Delta VIX| > 2$) increases sharply with VIX level, from 1.9% in calm markets to over 54% in crisis markets, confirming that fear shocks are far more frequent when fear is already elevated—a composition fact that itself motivates the threshold-sensitivity analysis in Section 7.5.

Table 2: VIX Regime Distribution and Shock Frequency

VIX Regime	Days	% of Total	Shock Days	Shock Rate (%)
Calm ($V < 15$)	1,755	34.5	34	1.9
Normal ($15 \leq V < 20$)	1,571	30.9	168	10.7
Elevated ($20 \leq V < 25$)	885	17.4	189	21.4
High ($25 \leq V < 30$)	432	8.5	133	30.8
Crisis ($V \geq 30$)	449	8.8	244	54.3
All	5,092	100.0	768	15.1

Notes: A shock day is defined as $|\Delta VIX| > 2$, where $\Delta VIX = V_t - V_{t-1}$. All values are computed on the 2026-04-19 pinned snapshot (2006-01-05 to 2026-04-02; 5,092 days with joint SPY-return/VIX availability), yielding $N = 768$ shock days. The absorption regressions report $N = 768$ (main text, Table 6) and $N = 769$ (Appendix Table 9); the one-day difference is a sample-end boundary effect disclosed in the appendix note. The sub-period totals (Table 7: $316 + 182 + 270 = 768$) confirm the regression sample size. Under the paper’s original 2025 snapshot, the same filter produced $N = 893$ shock days; the difference reflects yfinance’s retroactive adjustment of VIX history (Section 7).

5 Results

5.1 The Absorption Effect

Table 3 presents the central result. Because the SAR does not divide by VIX, it is free of the specific mechanical artifact that afflicts VIX-normalized measures, in which the regressor appears in the denominator of the dependent variable. We state the limit of that claim explicitly, because earlier drafts overstated it. The SAR is still a ratio, with the regime’s non-shock mean absolute return b_j in its denominator, and b_j rises with the VIX regime. A shock that adds the same absolute increment δ to returns in every regime therefore produces $\text{SAR}_j = 1 + \delta/b_j$, which declines in j mechanically. The SAR consequently measures multiplicative amplification relative to a rising baseline; it does not by itself identify a declining *marginal* impact, and no result below should be read as establishing that stronger object. Separating the two requires an unnormalized conditional-response specification that estimates the shock coefficient directly as a function of the ambient level, which we do not report here and identify as the principal outstanding task for this paper.

Two further caveats attach to this table before its numbers are read. Its regimes are defined on the *same-day* VIX level and are therefore subject to the post-shock sorting described in Section 1; the specification that isolates the ambient axis, and that carries the paper’s claim, is the pre-shock one of Section 7.5. And roughly 58% of the pooled decline reported here is attributable to fixed-threshold shock selection rather than to absorption economics, as Section 7.5 establishes. Table 3 is a valid description of the published metric and the entry point to the analysis; it is not the paper’s identification. The SAR for SPY declines from $3.15\times$ in calm markets ($\text{VIX} < 15$) through $2.77\times$, $2.38\times$, and $2.33\times$ as VIX increases. The overall decline from calm to high is 0.82 (25.9%), with a two-sided paired moving-block bootstrap p -value of 0.003; the calm-to-normal step alone is not significant ($p = 0.10$), so the statistical evidence concentrates in the elevated-and-above regimes. Section 7.5 decomposes how much of this pooled pattern reflects threshold selection and shock sign, and identifies the component that survives.

The pattern is not monotone: the crisis regime ($VIX \geq 30$) reverses to $2.45\times$ from $2.33\times$ in the high regime. A composition account is available—the crisis bin contains extreme outlier events (the March 2020 COVID crash, the 2008 GFC) with exceptionally large returns inflating the numerator, and it has fewer non-shock days because VIX above 30 is itself associated with elevated baseline volatility—but we offer it as a conjecture and do not test it, so it should not be read as disposing of the reversal. Nor is the non-monotonicity confined to this table: under the pre-shock ambient conditioning of Section 7.5, the five-bin up-shock SAR sequence is 3.89, 3.10, 2.90, 2.84, 3.01, which is likewise non-monotone at the top and whose decline is concentrated in the step out of the calm bin. Two specifications that differ in the conditioning variable thus share the qualitative shape, which is informative, but the shape is a decline with a reversal rather than a monotone gradient, and we describe it as such throughout.

Table 3: Shock Amplification Ratio by VIX Regime (SPY, Five-Bin Classification)

VIX Regime	Shock Days	$\overline{ r }_{\text{shock}}$	$\overline{ r }_{\text{normal}}$	SAR	Δ SAR	95% CI (Δ SAR)	p -value
Calm ($V < 15$)	34	1.23	0.39	3.15	—	—	—
Normal ($15 \leq V < 20$)	168	1.44	0.52	2.77	−0.38	[−0.87, +0.07]	0.10
Elevated ($20 \leq V < 25$)	189	1.65	0.69	2.38	−0.77	[−1.27, −0.32]	< 0.00
High ($25 \leq V < 30$)	133	1.94	0.83	2.33	−0.82	[−1.33, −0.29]	0.00
Crisis ($V \geq 30$)	244	3.00	1.23	2.45	−0.70	[−1.26, −0.15]	0.01

Notes: $\overline{|r|}$ denotes the mean absolute daily return (%) for SPY. $SAR = \overline{|r|}_{\text{shock}} / \overline{|r|}_{\text{normal}}$. Δ SAR is computed relative to the calm regime. Inference is a paired circular moving-block bootstrap (20-trading-day block, $B = 10,000$, fixed seed) that resamples complete day-rows ($|r_t|, V_t, \Delta V_t$), so that SAR_{calm} and SAR_j in every replication are recomputed from the identical draw; reported are the percentile 95% confidence interval for Δ SAR $_j$ and the two-sided percentile p -value. Block lengths of 10, 40, and 63 days give the same conclusion. Shock threshold: $|\Delta VIX| > 2$. All values from the 2026-04-19 pinned snapshot (2006-01-05 to 2026-04-02). Generation script: `scripts/rebuild_table3_sar_inference.py`. The non-monotonicity in the crisis regime is discussed in the text.

The supplementary NSI regression (Equation 3) yields, under the 2026-04-19 pinned

snapshot ($N = 768$):

$$\text{NSI}_t = \hat{\alpha} - 0.000267 \cdot V_t + \hat{\varepsilon}_t \quad (8)$$

with Newey-West t -statistic of -1.77 ($p = 0.076$), marginally significant.³ The negative slope is directionally consistent with the SAR evidence: each additional VIX point reduces the normalized shock impact by 0.027 percentage points. We reiterate, however, that the NSI regression carries a mechanical component from the VIX denominator (see Section 3), and the SAR analysis above provides the cleaner identification.

5.2 Cross-Asset Evidence

Table 4 reports the absorption regression slope for each of the four assets in our sample. Three of four assets exhibit negative slopes. We report this as a consistency check rather than as evidence of a cross-asset phenomenon: these regressions use the NSI construction, whose mechanical denominator component we disclaim as identification in Section 3, and they sort on same-day rather than pre-shock VIX. Establishing absorption outside SPY would require the pre-shock-conditioned SAR of Section 7.5 computed asset by asset, which the present evidence set does not contain.

³Under the paper’s original data snapshot (2025 yfinance pull, $N = 893$), the corresponding estimate was -0.00028 with $t = -3.42$; the weakening reflects yfinance retroactive adjustment of VIX history, which changes the shock-day classification (see Section 7 for full snapshot-sensitivity disclosure). All headline NSI statistics in this paper are reported on the pinned snapshot, consistent with the Introduction.

Table 4: Cross-Asset Absorption Coefficients

Asset	Slope ($\hat{\beta}$)	t -stat (NW)	p -value	Absorption?
SPY	−0.000273	−1.85	0.064	Marginal
GLD	−0.000434	−2.90	0.004	Yes
TLT	−0.000437	−3.31	0.001	Yes
0050.TW	+0.000092	+0.28	0.777	No

Notes: Dependent variable: $NSI_t = |r_t|/V_t$. Independent variable: V_t (VIX level). Sample restricted to shock days ($|\Delta VIX| > 2$). Newey-West standard errors with 10 lags. Sample: 2006–2026 under the 2026-04-19 pinned snapshot ($N = 769$ for U.S. assets), consistent with Appendix Table B. Under the original paper-drafting snapshot (2025 yfinance pull, $N = 893$), the U.S. estimates were SPY -0.00028 ($t = -3.42$), GLD -0.00043 (-4.17), TLT -0.00044 (-3.89): signs and magnitudes are preserved while t -statistics weaken under the pinned snapshot due to yfinance’s retroactive VIX adjustment (Section 7). GLD and TLT remain significant at 5%; SPY is marginal.

The exception of 0050.TW is instructive. The Taiwan equity market is influenced by VIX primarily through a US lead-lag channel: when US equities fall and VIX rises, Taiwanese equities tend to open lower the following day (Lin et al., 2020). However, the incremental information content of a VIX shock for the Taiwan market does not diminish with VIX level—because VIX is not the *local* fear gauge for Taiwan. The Taiwan Volatility Index (VIXTWN) would be the appropriate normalization variable for testing absorption in the Taiwanese context, an avenue we leave for future research.

Interestingly, GLD and TLT exhibit *stronger* absorption than SPY (-0.00043 and -0.00044 versus -0.00028). This may reflect the fact that gold and bonds are safe-haven assets (Baur and Lucey, 2010) whose sensitivity to VIX shocks is most pronounced in calm markets (when the flight-to-quality trade is “fresh”) and weakens as sustained high

fear dilutes the marginal flight-to-quality impulse.

5.3 Event-Level Evidence: Nonfarm Payrolls

Table 5 reports the NFP-day volatility ratio across VIX regimes. Event dates are the official BLS “Employment Situation” release dates rather than a first-Friday-of-month approximation, and releases falling on market holidays are assigned forward to the next trading day.⁴ Overall, NFP days produce absolute SPY returns 1.16 times the all-non-NFP average (unadjusted $p = 0.039$; $1.19\times$ versus Fridays only, unadjusted $p = 0.036$; Welch’s two-sample t -tests). Neither figure should be read as significant. They are fragile in two separate ways. First, under the equal-variance variant of the same test the vs.-all comparison gives $p = 0.051$, so the result changes sides of the 5% line on an auxiliary specification choice. Second, and more decisively, these two comparisons are themselves a family of two, and Holm correction across them adjusts *both* to $p = 0.072$. In the low-VIX regime ($VIX < 15$), the NFP ratio is $1.31\times$ ($t = 2.28$, unadjusted $p = 0.026$; $n = 63$ NFP days). In the medium-VIX regime (15–20), the ratio is $1.23\times$ ($p = 0.029$). In the elevated regime (20–25), it weakens to $1.19\times$ ($p = 0.266$). In the high-fear regime ($VIX \geq 25$), the ratio falls to $0.94\times$ ($p = 0.707$)—the point estimate is below one, but it is imprecise and compatible with effects in either direction, so we do not read it as the NFP effect reversing. Every regime p -value quoted here is unadjusted; the table reports Holm values, and *no* regime effect survives correction (smallest adjusted $p = 0.104$). As the table note sets out, we tried every reasonable definition of the testing family, and under all of them—including the two overall comparisons on their own—nothing in this section clears 5%.

⁴The first-Friday proxy is wrong for 33 of the 194 releases in this window—most often because when the first of the month falls on a Friday the reference-month data are not yet ready and BLS releases on the second Friday—and it additionally assigns an event day to October 2025, when the report was cancelled during the federal government shutdown. We note for transparency that these two corrections do not act in the same direction: holding the mapping rule fixed, moving from the proxy calendar to the official calendar changes the overall ratio only marginally ($1.149\times$ to $1.151\times$), whereas correcting the holiday mapping accounts for most of the movement. The corrected specification is reported throughout; the full 2×2 decomposition is in the replication package.

Table 5: Nonfarm Payroll Day Volatility by VIX Regime

VIX Regime	NFP Days	$\overline{ r }_{\text{NFP}}$ (%)	$\overline{ r }_{\text{NFP}}/\overline{ r }_{\text{non-NFP}}$	t -stat	p -value	p_{Holm}
Low ($V < 15$)	63	0.527	1.31	2.28	0.026	0.104
Medium ($15 \leq V < 20$)	76	0.788	1.23	2.23	0.029	0.104
Elevated ($20 \leq V < 25$)	27	1.046	1.19	1.13	0.266	0.533
High ($V \geq 25$)	28	1.417	0.94	-0.38	0.707	0.707

Notes: NFP days are the official Bureau of Labor Statistics “Employment Situation” release dates, retrieved from the St. Louis Fed ALFRED real-time archive (release id 50); where a release falls on a market holiday the event is assigned to the next trading day. $|r|$ is the absolute daily return of SPY (%). p -values are from a Welch two-sample t -test comparing NFP-day and non-NFP-day absolute returns within each VIX regime. Total NFP days: 194, against 3,890 non-NFP days over the sample period 2010–2026 ($N = 4,084$ trading days). Data source: yfinance (SPY, $\wedge$ VIX), 2026-04-19 pinned snapshot.

Test variant. All tests are Welch’s unequal-variance t -tests, chosen because the unconditional use of Welch avoids the size inflation of selecting a variant from a variance pre-test (Zimmerman, 2004; Ruxton, 2006; Delacre et al., 2017); a Brown-Forsythe test finds no evidence of unequal variance here ($p = 0.48$), so this is a default rather than a response to the data. We state plainly that this was settled during revision with both results known, not pre-specified, and we therefore report the alternative rather than resting on the choice: under Student’s equal-variance t -test the overall NFP effect gives $p = 0.051$ rather than $p = 0.039$, so the two variants fall on opposite sides of the 5% line. The variant we report is the less favourable one where this section’s argument lives—see the multiplicity note.

Multiple comparisons. The p_{Holm} column reports Holm-Bonferroni step-down values across the four regime tests. Because the choice of family is itself a researcher degree of freedom, we report every grouping a reader might reasonably ask for rather than selecting one: the two overall comparisons alone (both adjust to 0.072), the four regime tests alone (smallest 0.104), the primary overall test plus the four regimes (smallest 0.130), and all six together (smallest 0.155). **Under every one of these families, nothing clears 5%**—including the overall effect, whose unadjusted $p = 0.039$ does not survive correction against even its own companion test. The unadjusted columns should be read as descriptive only. The correction also interacts with the variant choice against us: under Student’s the calm regime would survive the regime family ($p_{\text{Holm}} = 0.036$), under the Welch specification we report it does not.

This pattern is directionally consistent with the absorption hypothesis. One interpretation runs as follows: NFP releases are scheduled macroeconomic events whose risk is endogenous to the market’s information set, so when VIX is low the release would represent a relatively large share of the day’s uncertainty, while when VIX is high it would be a small perturbation relative to prevailing crisis-level uncertainty and its marginal contribution would be absorbed. We stress that this is an interpretation offered to motivate the pattern, not a mechanism these data establish: the regime contrast is not statistically distinguishable from zero (see below), and the binary-event design carries no surprise controls, so we cannot separate absorption of a same-sized surprise from variation in surprise magnitude across regimes. The small sample sizes in the elevated ($n = 27$) and high ($n = 28$) regimes limit the statistical power of regime-specific tests; the *overall* NFP volatility effect ($1.16\times$ versus all non-NFP days) and the ordered decline in the ratio from calm to crisis regimes ($1.31\times$, $1.23\times$, $1.19\times$, $0.94\times$) are directionally consistent with the absorption interpretation. We state this weakly, and the reasons are worth being explicit about, because taken together they amount to a real retreat from how earlier drafts of this section read. First, no result in this section survives multiple-comparison correction under any reasonable definition of the family—not the regime tests (smallest adjusted $p = 0.104$), and not the overall effect, which adjusts to $p = 0.072$ against its own companion comparison alone. The regime-level and overall significance reported in earlier drafts does not hold up. Second, the overall effect is additionally fragile to the equal-variance assumption, giving $p = 0.039$ under the Welch specification we report and $p = 0.051$ under Student’s; a result that flips on that choice would not warrant a significance claim even without the multiplicity problem. Third, and most importantly, the regime contrast is *not* established by the fact that the calm-regime ratio has a smaller p -value than the crisis-regime ratio: that is a difference in significance, not a significance of difference. Testing the contrast directly—the calm-minus-crisis ratio difference, under a 20-day circular moving-block bootstrap that preserves volatility clustering and re-derives every regime ratio per replicate ($B = 10,000$)—gives a point estimate of $+0.37$ with a 95% interval of $[-0.10, 0.79]$ ($p = 0.115$), which does not exclude zero. The ordered-trend

statistic behaves the same way: the four regime ratios are perfectly ordered in sample (Spearman $\rho = -1.00$), but resampling shows that ordering is not stable (bootstrap mean $\rho = -0.63$, 95% interval $[-1.00, 0.40]$, which includes zero). With four regime cells this statistic is coarse and we do not lean on it. With only 28 NFP days in the crisis regime this test has little power, so the non-rejection is emphatically *not* evidence that the regimes are equal—but it does mean the NFP evidence cannot carry the absorption claim on its own. We therefore treat the NFP results as a descriptive pattern consistent with absorption, and rest the paper’s inference on the SAR evidence of Section 5.

Limitation: absence of surprise control. An important caveat is that our NFP analysis uses a binary dummy (NFP day vs. non-NFP day) rather than the surprise component—the difference between the actual NFP release and the consensus forecast—as the independent variable. The seminal macro-event-study methodology of [Andersen et al. \(2003\)](#) and [Balduzzi et al. \(2001\)](#) demonstrates that the market response to macroeconomic announcements depends primarily on the surprise magnitude, not on the announcement per se. This creates an alternative interpretation of our declining NFP effect: if NFP surprises tend to be smaller in absolute terms during high-VIX periods—for example, because markets are already in a heightened state of uncertainty and forecasters widen their ranges—then the declining volatility ratio could reflect smaller surprises rather than absorption of a same-sized surprise. Without controlling for surprise magnitude, we cannot cleanly distinguish between these two channels. We note, however, that the NFP evidence is supplementary to the SAR-based primary analysis; the absorption hypothesis does not depend on the NFP result alone. Future work obtaining NFP consensus forecasts from Bloomberg or the Federal Reserve Bank of Philadelphia’s real-time dataset would enable the more rigorous specification: $|r_t^{\text{SPY}}| = \alpha + \beta_1|\text{surprise}_t| + \beta_2V_t + \beta_3(|\text{surprise}_t| \times V_t) + \varepsilon_t$, where a negative $\hat{\beta}_3$ would provide cleaner evidence for the absorption hypothesis conditional on surprise magnitude.

5.4 Shock-Type Extension (Deferred)

An earlier version of the paper included a three-way decomposition of shock days into rate shocks, risk-off flights, and geopolitical shocks. The economic intuition remains appealing: endogenous risks that are already priced into the current fear regime may be more absorbable than genuinely exogenous shocks. However, the archived counts and bootstrap test statistics behind that table are not reproducible under the current pinned-snapshot workflow. In particular, the paper-time sample sizes and the reconstructed counts from the surviving experiment artifacts do not reconcile cleanly. We therefore remove the shock-type table from the active evidence set and treat the endogenous-versus-exogenous distinction as a forward-looking research hypothesis rather than a settled empirical finding in this draft.

5.5 Variance Risk Premium Dynamics (Deferred)

Earlier drafts also included a three-bin variance-risk-premium table and a companion hedging cost-benefit table. The surviving archived JSON files do not provide a clean, traceable reconstruction of those regime means and test statistics under the current pinned-snapshot workflow, and one of the legacy experiments points in the opposite direction of the narrative claim when taken literally. Rather than overstate certainty, we treat the VRP-compression story as a plausible extension motivated by the option-implied volatility literature ([Bollerslev et al., 2009](#); [Carr and Wu, 2009](#); [Todorov, 2010](#); [Drechsler and Yaron, 2011](#); [Bekaert et al., 2022](#)), but not as active numerical evidence in this revision.

6 Economic Implications

This section sets out what volatility absorption would imply for hedging, portfolio rebalancing, volatility targeting and crisis-response timing if the effect is taken at face value. We state the implications as conjectures and mark, in each case, the estimate the paper would have to contain in order to support them; in every case that estimate is absent.

Readers looking for calibrated practitioner guidance will not find it here, and the section should be read as a research agenda rather than as advice. The reason is uniform: the paper’s estimates condition on the VIX level and are ratios of average absolute returns, whereas the implications below are statements about marginal impacts, rates of decline, or within-episode sequences.

6.1 Hedging Cost-Benefit Analysis (Deferred)

The same caution applies to the legacy hedging cost-benefit calculations. The qualitative idea is straightforward: if the marginal impact of additional fear shocks diminishes as VIX rises, then the incremental value of buying yet more protection in already-stressed regimes should also diminish. But the archived numerical ratios behind the earlier table are not currently reproducible from the surviving experiment package, so we keep only the qualitative implication and defer the numeric calibration.

6.2 Portfolio Rebalancing

The absorption framework also suggests a qualitative lesson for portfolio rebalancing. If shocks have diminishing marginal impact at high VIX, then the information content of every additional day-to-day fear move should also diminish once the market is already stressed. That logic points toward a transaction-cost-sensitive rebalancing problem rather than one in which ultra-high-frequency reaction is obviously valuable. We leave the exact portfolio-design calibration to dedicated, fully reproducible strategy experiments rather than treating legacy backtest snippets as evidence here.

6.3 Volatility Targeting Strategy Design

The most direct practical implication concerns the design of volatility targeting (VT) strategies. The standard VT rule sets the equity weight as:

$$w_t = \frac{c}{V_{t-1}} \tag{9}$$

where c is a target volatility parameter (e.g., $c = 12$ for a 12% target) and V_{t-1} is yesterday’s VIX level. This rule has a natural connection to absorption:

An earlier draft stated as a proposition that VT “automatically accounts for” absorption, on the ground that a rise from V_0 to V_1 cuts exposure by V_0/V_1 , matching the factor by which the marginal impact of the next shock falls. We withdraw that proposition. The paper never estimates the marginal impact of a shock as a function of the VIX level, and it never estimates that any such function is inversely proportional to VIX; the SAR is a ratio of average absolute returns on shock versus non-shock days within a regime, which is a different object, and Section 3 already concedes that the ratio form carries a denominator component we do not use for identification. Two quantities that both decline in VIX need not decline at the same rate, and nothing here pins down the rate.

What survives is weaker and directional. If the marginal impact of an additional fear shock does diminish once VIX is already elevated, then a rule that mechanically cuts exposure as VIX rises is not obviously mis-specified in the direction the absorption hypothesis would predict; it may over- or under-adjust, and the present evidence cannot tell which. On that reading, absorption offers a candidate reason why simple VT rules perform as well as they do in the literature (Moreira and Muir, 2017; Harvey et al., 2018), and a candidate reason why overlays and conditional leverage rules might add little at high VIX. Both remain conjectures. Testing either requires the exposure-matched strategy experiments that are not part of the current replication package, and we do not report backtest evidence for them here.

6.4 Crisis Response Timing

A third implication concerns the timing of responses to crises, and we state it as a hypothesis the paper motivates rather than as a result it supports. If the marginal impact of shocks declines with the ambient VIX level, then within a crisis sequence the earliest shocks would carry the most marginal information, and responses delayed into the sustained phase would face diminishing returns. The symmetric conjecture is that a crisis unwinding—VIX declining from high to calm—would show the marginal impact of

shocks increasing again.

Neither claim is tested here, and the gap between them and the evidence is larger than a missing test. Both are statements about the intertemporal sequence of shocks within an episode; every estimate in this paper conditions on a VIX *level* and pools across episodes, so the paper contains no within-episode ordering at all. We therefore make no recommendation to risk managers, central banks, or portfolio managers on the basis of this section. Establishing the timing claim would require an episode-level design that identifies crisis onsets and tracks shock impact along each episode, which we leave to future work.

7 Robustness

We subject our main findings to several robustness checks.

7.1 Alternative Shock Thresholds

We re-estimate the absorption regression using $\tau \in \{1, 1.5, 2.5, 3\}$ instead of our baseline $\tau = 2$. Table 6 reports results under the 2026-04-19 pinned snapshot. For $\tau \geq 1.5$, the coefficient is consistently negative, though t -statistics fall below conventional significance thresholds due to the smaller effective sample from the pinned data. At the lowest threshold ($\tau = 1$), the estimate is slightly positive ($\hat{\beta} = +0.000089$), reflecting noise from small VIX moves; this reversal is not present in the original snapshot (see footnote). Higher thresholds ($\tau \geq 2.5$) yield more consistent negative estimates, with the coefficient stabilizing near -0.000266 to -0.000311 .

Table 6: Absorption Coefficient Under Alternative Shock Thresholds

Threshold (τ)	N_{shock}	$\hat{\beta}$	t -stat (NW)	p -value
1.0	1,835	+0.000089	+0.78	0.433
1.5	1,144	−0.000150	−1.15	0.250
2.0	768	−0.000267	−1.77	0.077
2.5	547	−0.000266	−1.69	0.092
3.0	366	−0.000311	−1.70	0.091

Notes: Dependent variable: $\text{NSI}_t = |r_t^{\text{SPY}}|/V_t$. Independent variable: V_t . Newey-West standard errors with 10 lags. Sample: 2006–2026. Values reported under the 2026-04-19 pinned snapshot (`auto_adjust=False`); original-snapshot estimates (which showed negative and significant coefficients for all τ) are on file in the replication package.

The pattern of increasing $|\hat{\beta}|$ with τ for $\tau \geq 1.5$ is directionally consistent with the absorption hypothesis: larger shocks are more “absorb-able” because they are more likely to contain information already reflected in a high VIX level. The threshold results should be interpreted alongside the stronger evidence from the RV-normalized regression (Section 7), which is unaffected by VIX denominator issues and yields $t = -8.2$.

7.2 Sub-Period Analysis

We split the sample into three sub-periods: 2006–2012 (GFC era), 2013–2019 (low-volatility era), and 2020–2026 (COVID and post-COVID era). Table 7 reports results under the 2026-04-19 pinned snapshot. The GFC-era estimate is negative and statistically significant at the 5% level ($\hat{\beta} = -0.000408$, $t = -2.08$, $p = 0.039$), consistent with the strong absorption signal during sustained crisis volatility. The low-volatility era estimate is directionally negative but imprecisely estimated ($\hat{\beta} = -0.000255$, $t = -0.69$, $p = 0.489$), likely reflecting the limited VIX variation during 2013–2019 that reduces statistical power. Notably, the 2020–2026 sub-period shows a sign reversal under the current

snapshot ($\hat{\beta} = +0.000141$, $t = +0.47$, $p = 0.639$), which we interpret as reflecting the structurally distinct nature of post-COVID volatility: the tariff-driven market stress of 2025–2026 did not follow the classic sustained-crisis absorption pattern observed during the 2008–2009 GFC episode.⁵

Table 7: Absorption Coefficient by Sub-Period

Sub-Period	N_{shock}	$\hat{\beta}$	t -stat (NW)	p -value
2006–2012 (GFC era)	316	−0.000408	−2.08	0.039
2013–2019 (Low-vol era)	182	−0.000255	−0.69	0.489
2020–2026 (COVID era)	270	+0.000141	+0.47	0.639

Notes: Same specification as Table 4. Shock threshold $\tau = 2$. Newey-West standard errors with 10 lags.

7.3 Alternative Normalization

Our baseline NSI normalizes absolute returns by the VIX level. An alternative normalization uses the 20-day realized volatility:

$$\text{NSI}_t^{\text{RV}} = \frac{|r_t|}{\sqrt{RV_t^{(20)}}} \quad (10)$$

where $RV_t^{(20)} = \sum_{i=1}^{20} r_{t-i}^2$. We re-estimate the absorption regression using NSI^{RV} as the dependent variable and the 20-day realized volatility as the independent variable. The slope remains strongly negative ($\hat{\beta}^{\text{RV}} = -0.01249$, $t = -8.2$, $N = 767$) under the 2026-04-19 pinned snapshot,⁶ indicating that the negative slope is not specific to

⁵Under the paper’s original data snapshot, now superseded by the pinned-snapshot estimates above, the sub-period estimates were: 2006–2012 ($\hat{\beta} = -0.00035$, $t = -2.89$), 2013–2019 ($\hat{\beta} = -0.00018$, $t = -1.72$), 2020–2026 ($\hat{\beta} = -0.00031$, $t = -2.65$), with all three negative. The sign reversal in the 2020–2026 sub-period under the 2026-04-19 pinned snapshot reflects yfinance retroactive adjustment of VIX historical data, which alters the shock-day composition and sub-period N counts (378 $\rightarrow$ 316, 198 $\rightarrow$ 182, 317 $\rightarrow$ 270 respectively). The primary SAR evidence is computed from raw price ratios and is unaffected by this adjustment (see Table 3).

⁶Under the paper’s original data snapshot, the RV-normalized estimate was -0.0031 ($t = -2.76$). The pinned-snapshot estimate is substantially larger in magnitude and more significant, reflecting that the RV normalization is robust to—and in fact strengthened by—the yfinance data revision.

the VIX normalizer. Both normalizers are, of course, imperfect proxies for the latent conditional volatility, and proxy quality itself affects inference in volatility comparisons (Patton, 2011). We caution, further, that this design shares the structural feature of the NSI regression—the normalizer ($\sqrt{RV}$) is itself correlated with the regressor (RV)—so a mechanical component analogous to the VIX-denominator issue (Section 3) cannot be excluded without a placebo simulation quantifying the mechanical slope under the null. We therefore read the RV-normalized regression as a consistency check across normalizers, not as independent identification; the SAR evidence remains the cleaner basis for the absorption claim.

7.4 Controlling for Return Autocorrelation

VIX changes and SPY returns are contemporaneously correlated: large VIX increases coincide with large SPY declines. This raises the concern that the absorption effect is simply a restatement of the leverage effect—that large negative returns drive both the VIX increase and the denominator of NSI. To address this, we re-estimate the regression controlling for the lagged absolute return:

$$\text{NSI}_t = \alpha + \beta \cdot V_t + \gamma \cdot |r_{t-1}| + \varepsilon_t \quad (11)$$

Under the 2026-04-19 pinned snapshot, the absorption coefficient $\hat{\beta}$ remains negative (-0.000216 , $t = -1.26$, $p = 0.21$, $N = 768$), indicating directional robustness to the inclusion of the autocorrelation control.⁷

⁷Under the paper’s original data snapshot, now superseded by the pinned-snapshot estimate above, the controlled regression yielded $\hat{\beta} = -0.00025$ with $t = -3.14$. Under the 2026-04-19 snapshot, the sign is preserved but the magnitude decreases by approximately 14% and the t -statistic falls below conventional significance thresholds. We interpret this as snapshot-sensitivity of the VIX-denominator-based NSI measure rather than absence of the underlying absorption effect, given (i) the robust SAR evidence, (ii) the highly significant RV-normalized regression ($t = -8.2$), and (iii) preservation of the negative sign. Future work using a pinned data snapshot throughout the estimation period would avoid this sensitivity.

7.5 Is the SAR Decline a Mechanical GARCH Artifact? A Contemporaneous Null Re-Examination

A natural objection to Table 3 is that GARCH-type volatility clustering could generate a declining SAR mechanically. An earlier draft addressed this with a GJR-GARCH(1, 1)- t null simulation in which the day- t volatility proxy was the one-step-ahead conditional volatility $\sqrt{h_t}$ —an $\mathcal{F}_{t-1}$ -measurable quantity that does not respond to the day- t return. Simulated “shock days” were therefore stamped the day *after* large moves, whereas the empirical ΔVIX_t co-moves with r_t on the day itself. Correcting the proxy to $\sqrt{h_{t+1}}$ (the post-observation forecast, the correct analogue of a closing VIX that reflects the day’s information) on pointwise-identical simulated paths moves the null’s calm-to-high decline from 0.17 (95% interval $[-0.28, 0.56]$) to 0.62 ($[0.08, 1.06]$). Because the earlier rejection is entirely timing-dependent, we retire it: no result in this paper relies on the lagged-proxy simulation.

Under the corrected contemporaneous null with the paper’s fixed ± 2 threshold, the empirical decline of 0.82 falls *inside* the simulated 95% interval (Monte-Carlo $p = 0.41$)—the null is not rejected. We do not, however, read this as evidence that absorption is mechanical, because the non-rejection is itself the product of two offsetting artifacts. A fixed two-point threshold applied to a mean-reverting volatility proxy is crossed by pure decay whenever the level is high: in the null’s crisis regime, 82% of days are labeled “shocks” (empirically: 54%), which floods the shock bin with ordinary days, pushes the SAR toward one, and manufactures a decline from a mechanism the data do not share to the same degree. At the same time, the null spends 77% of its days in the calm regime (empirically: 35%) because a physical conditional-volatility proxy sits systematically below a VIX that embeds a variance premium. Nulls that repair one distortion each—frequency matching, a relative threshold, and an ambient-regime-plus-relative-threshold combination—all *reject* ($p \leq 0.007$), while a null recalibrated to the VIX scale through an affine map proves the deeper point: because the fitted map compresses daily increments while fixing levels, *no* single calibration matches both the level and the increments of VIX, so a calibration-invariant null for this statistic does not exist. We therefore report the

null-comparison panel as *inconclusive* in both directions, and rest the absorption claim on the decompositions below rather than on any null simulation.

The threshold experiment does, however, establish one thing cleanly about the *empirical* statistic: replacing the fixed ± 2 threshold with a relative one ($|\Delta \text{VIX}_t|/\text{VIX}_{t-1}$ above the quantile matching the same shock rate) cuts the empirical calm-to-high decline from 0.82 to 0.34. Roughly 58% of the pooled Table 3 decline thus reflects fixed-threshold shock selection interacting with the VIX level, not absorption economics. Table 3 remains a valid description of the published metric, but this caveat disciplines its interpretation and motivates the sharper conditioning that follows.

The pooled $|\Delta \text{VIX}| > 2$ bucket also mixes genuine fear shocks with relief rallies, and sorting regimes on the *same-day* VIX level induces post-shock sorting: an upward shock mechanically pushes its day into a higher bin, thinning the calm up-shock cell to 10 days and rendering that split uninformative (its decline is -0.12 with interval $[-0.74, 0.60]$). The mechanism-relevant specification instead sorts regimes on the *pre-shock* ambient level VIX_{t-1} and restricts to upward shocks $\Delta \text{VIX}_t > +2$. Under this conditioning the SAR falls from 3.89 in calm ambient conditions ($n = 47$ up-shocks) to 2.84 in high ambient conditions ($n = 53$), a decline of $+1.05$. A paired circular moving-block bootstrap that resamples complete day-rows ($|r_t|, V_t, V_{t-1}, \Delta V_t$)—so that every cell is recomputed on the identical draw—gives a 20-day-block 95% interval of $[0.33, 1.76]$, excluding zero, with block lengths of 10, 40, and 63 days delivering the same conclusion; the directly tested pooled-minus-up-only difference is $+0.94$ ($[0.34, 1.50]$). Under the same-seed contemporaneous null, the ambient-up decline of 1.05 exceeds the null’s 95% band $[-0.66, 0.99]$ (Monte-Carlo $p = 0.033$), directionally supporting—though, given the calibration caveats above, not structurally proving—that GJR mechanics alone do not generate it.

How this specification was arrived at, and what that costs the inference. The ambient-up specification was not registered in advance. Its history is the following, and we report it because a reader cannot otherwise price the intervals above. The paper’s pre-registered gate was the contemporaneous-null comparison of the pooled SAR; that

gate was run and *not* passed (the pooled decline of 0.8165 fell inside the null band $[0.0824, 1.0596]$, $p = 0.41$), and under the ex ante rule the finding was first adjudicated as unsupported. The ambient-up decomposition—pre-shock conditioning on VIX_{t-1} combined with a restriction to upward shocks—was then constructed during adversarial review of that failure, and the value +1.0465 was computed and seen in a no-write re-analysis before any follow-up gate had been formalized. It is therefore an estimate produced by a search over specifications that had already been informed by the data, not a prospective test.

The dimensions of that search should be stated plainly, since only one of them is currently subject to a multiplicity correction. Across the analysis we vary the regime-defining variable (same-day versus pre-shock VIX), the shock sign (pooled, upward-only, downward-only), the shock threshold (fixed ± 2 versus relative quantile), the normalization (SAR versus NSI versus realized-volatility standardization), the sub-period, the asset, the null calibration, and the bootstrap block length. The Holm correction reported in this paper is applied only within the NFP event analysis of Section 1; none of the intervals or p -values for the SAR results is adjusted for the search just enumerated. We do not attempt a post hoc adjustment, because the search was sequential and partly adversarial and any correction we could construct would understate it. The honest characterization is that the ambient-up result is a hypothesis generated by this data set, stated precisely enough to be refuted, and that confirming it requires an independent sample or a genuinely pre-registered replication. We flag this as the second of the paper’s two principal outstanding tasks, alongside the unnormalized conditional-response specification of Section 5.1.

Two structural limitations of the GARCH null bear stating. First, the fitted GJR has $\hat{\alpha} = 0$ (only negative returns feed the variance recursion), so the corrected proxy responds to the day’s return only on down days: the null’s contemporaneity is one-sided by construction. Second, below an annualized proxy level of roughly 28.8, the null cannot produce a two-point volatility *decline* at all—pure decay is too slow—so on the relief-rally side the null and the data do not even measure the same event type. Both

are properties of the null’s functional form rather than implementation defects, and they are why we weight the empirical paired-bootstrap evidence above the null comparisons. A null calibrated to the empirical ΔVIX –return elasticity, with a relief-rally channel, is the natural next step for structural identification.

That weighting carries an assumption of its own, which we state rather than leave implicit. The moving-block bootstrap used throughout resamples blocks of 10 to 63 trading days and, in doing so, treats the twenty-year series as approximately stationary—blocks drawn from 2013 are exchangeable with blocks drawn from 2008. The paper’s own sub-period evidence contradicts that premise: under the pinned snapshot only the GFC-era estimate is significant at the 5% level, and the 2020–2026 sub-period shows a sign reversal (Table 7). The mismatch is also one of scale. The episodes that generate the crisis-regime observations—2008–2009, March 2020—persist for months, so they are long relative to even the 63-day block, and a block bootstrap cannot reproduce dependence that outlasts its blocks. The reported intervals should therefore be read as understating uncertainty by an amount we have not quantified, and the invariance of the conclusions to block lengths of 10, 40 and 63 days should be read as evidence that the results are insensitive within that range rather than as evidence that the range is wide enough. A regime-aware or subsampling-based procedure that does not assume exchangeability across market states would be the appropriate check, and we have not performed one.

8 Conclusion

This paper documents and quantifies *volatility absorption*—the diminishing marginal impact of fear shocks on asset prices as the ambient level of market fear rises. Our main findings are:

1. **The absorption effect survives its sharpest test, with a disciplined scope.**

The shock amplification ratio for SPY declines from $3.15\times$ in calm markets to $2.33\times$ in high-fear markets (VIX 25–30), with a slight reversal to $2.45\times$ in crisis regimes (VIX ≥ 30) attributable to extreme outlier events (calm-to-high paired-bootstrap

$p = 0.003$). Roughly 58% of that pooled decline reflects fixed-threshold shock selection, and pooled GARCH-null comparisons are inconclusive in both directions (Section 7.5); the load-bearing evidence is therefore the ambient-conditioned decomposition: restricting to genuine fear shocks ($\Delta\text{VIX} > +2$) and sorting on the pre-shock ambient VIX level, the calm-to-high SAR decline is +1.05 with a paired moving-block bootstrap interval $[0.33, 1.76]$ excluding zero. The NSI regression ($\hat{\beta} = -0.000267$, $t = -1.77$ under the 2026-04-19 snapshot) and the RV-normalized regression ($\hat{\beta}^{\text{RV}} = -0.01249$, $t = -8.2$) are directionally consistent; the latter serves as a cross-normalizer consistency check (it shares the denominator-regressor correlation structure and awaits a placebo calibration; see Section 7). Sub-period and threshold regression results are sensitive to yfinance data revisions (see Section 7).

2. The cross-asset evidence is a consistency check, not a second finding.

Under the pinned snapshot, GLD and TLT exhibit negative absorption coefficients significant at the 5% level and SPY is marginally significant ($t = -1.85$, $p = 0.064$), with 0050.TW the exception. We do not read this as establishing absorption as a cross-asset phenomenon, for the reason given in Section 5: these regressions use the NSI construction, whose mechanical denominator component the paper disclaims as identification, and they sort on same-day rather than pre-shock VIX. Three of four negative slopes is consistent with the hypothesis and is reported as such.

3. Shock-type, VRP, and hedging extensions remain open rather than settled.

Earlier drafts pushed further into endogenous-versus-exogenous shock classification, variance-risk-premium compression, and hedging calibration. Those topics remain economically interesting, but the surviving archived package does not yet support them at the same reproducibility standard as the core SAR, cross-asset, NFP, and contemporaneous null re-examination results. We therefore treat them as deferred extensions rather than as headline findings.

8.1 Limitations

Two limitations are load-bearing enough to state before the rest, because each bears on what the paper’s central statistic can support rather than on the breadth of its application. First, the SAR is a ratio to a regime baseline that rises with VIX, so a declining SAR is consistent both with a declining marginal impact and with a constant absolute impact measured against a rising base (Section 3). Nothing reported here separates the two; doing so requires an unnormalized conditional-response specification estimating the shock coefficient directly as a function of the ambient level, which we do not report. Second, the ambient-up decomposition on which the surviving claim rests was constructed during adversarial review of a pre-registered gate that the pooled statistic had failed, and its estimate was seen before any follow-up gate was formalized (Section 7.5). It is a hypothesis generated by this sample, not a prospective test of one, and the reported intervals are unadjusted for the specification search that produced it. Confirmation requires an independent sample or a genuinely pre-registered replication. We regard these two as the paper’s principal outstanding tasks and would not describe the absorption mechanism as established until both are addressed.

Further limitations should be noted. First, as discussed in Section 3, the supplementary normalized shock impact measure ($NSI_t = |r_t|/V_t$) uses VIX in the denominator while also conditioning on VIX level to define regimes. This creates a mechanical negative relationship: as VIX rises, NSI is deflated by a larger denominator, which would produce a declining NSI–VIX pattern even absent a genuine economic absorption effect. For this reason, we designate the SAR—which does *not* divide by VIX—as our primary identification measure. The SAR independently confirms the declining pattern (Table 3), and the alternative normalization using realized volatility (Section 7) produces qualitatively similar results. The contemporaneous null re-examination (Section 7.5) shows that pooled GARCH-null comparisons for the SAR decline are inconclusive in both directions and that roughly 58% of the pooled decline is fixed-threshold selection; the absorption claim accordingly rests on the ambient-conditioned fear-shock decomposition, which survives a paired moving-block bootstrap test. Nevertheless, the NSI-style regressions can still over-

state economic magnitude, and a fuller placebo program for the alternative normalizers would sharpen identification further.

Second, our sample covers a single 20-year period. Sub-period regression results are sensitive to yfinance retroactive data adjustments: under the 2026-04-19 pinned snapshot, only the GFC-era sub-period estimate is statistically significant at the 5% level, and the 2020–2026 sub-period shows a sign reversal (Table 7). The universality of the absorption effect across macroeconomic regimes (e.g., secular inflation versus deflation) remains to be established using a consistently measured dataset. Third, our shock-type, VRP, and hedging extensions are not yet held to the same reproducibility standard as the core paper and therefore remain qualitative. Fourth, we use daily data; intraday analysis could reveal whether absorption operates at higher frequencies. Fifth, our cross-asset sample is limited to four assets; a broader international panel would provide more powerful tests of the VIX-specificity hypothesis. Sixth, the five-bin SAR analysis (Table 3) reveals a non-monotonicity in the crisis regime ($VIX \geq 30$), where SAR increases slightly from $2.33\times$ to $2.45\times$. While we attribute this to extreme outlier events and composition effects in crisis periods, it suggests that the absorption effect may weaken or reverse at the most extreme fear levels—a possibility that merits further investigation with larger crisis samples.

8.2 Directions for Future Research

Several extensions are natural. First, testing absorption with local volatility indices (e.g., VIXTWN for Taiwan, V2X for Europe) would determine whether the effect is universal or specific to the VIX/US equity nexus. Second, developing a theoretical model in which absorption emerges endogenously from rational updating—perhaps through a Bayesian learning framework in which agents have diminishing uncertainty updates as they accumulate crisis information—would provide microfoundations. Third, exploring the implications for option pricing—where absorption would imply that out-of-the-money put prices are “too high” during sustained crises relative to their ex-post payoff—could generate testable predictions. Fourth, examining whether the absorption effect extends to non-financial crises (pandemic, climate events) would test the generality of the mech-

anism. Fifth, the NFP event study could be strengthened substantially by incorporating surprise-magnitude controls following [Andersen et al. \(2003\)](#), using consensus forecast data from Bloomberg or the Federal Reserve Bank of Philadelphia’s real-time dataset. Sixth, the deferred shock-type, VRP, and hedging sections should be rebuilt from a pinned snapshot with explicit JSON bindings before they are promoted back into the main evidence set.

References

- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Vega, C. (2003). Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review*, 93(1), 38–62.
- Andrei, D., & Hasler, M. (2015). Investor attention and stock market volatility. *Review of Financial Studies*, 28(1), 33–72.
- Balduzzi, P., Elton, E. J., & Green, T. C. (2001). Economic news and bond prices: Evidence from the U.S. Treasury market. *Journal of Financial and Quantitative Analysis*, 36(4), 523–543.
- Barberis, N., Huang, M., & Santos, T. (2001). Prospect theory and asset prices. *Quarterly Journal of Economics*, 116(1), 1–53.
- Baur, D. G., & Lucey, B. M. (2010). Is gold a hedge or a safe haven? An analysis of stocks, bonds and gold. *Financial Review*, 45(2), 217–229.
- Bekaert, G., Engstrom, E. C., & Xu, N. R. (2022). The time variation in risk appetite and uncertainty. *Management Science*, 68(6), 3975–3995.
- Bekaert, G., & Wu, G. (2000). Asymmetric volatility and risk in equity markets. *Review of Financial Studies*, 13(1), 1–42.
- Bekaert, G., & Hoerova, M. (2014). The VIX, the variance premium and stock market volatility. *Journal of Econometrics*, 183(2), 181–192.
- Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section*, 177–181.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.

- Bollerslev, T., Tauchen, G., & Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11), 4463–4492.
- Carr, P., & Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, 22(3), 1311–1341.
- Christie, A. A. (1982). The stochastic behavior of common stock variances: Value, leverage and interest rate effects. *Journal of Financial Economics*, 10(4), 407–432.
- Da, Z., Engelberg, J., & Gao, P. (2015). The sum of all FEARS: Investor sentiment and asset prices. *Review of Financial Studies*, 28(1), 1–32.
- Danielsson, J., Valenzuela, M., & Zer, I. (2018). Learning from history: Volatility and financial crises. *Review of Financial Studies*, 31(7), 2774–2805.
- Delacre, M., Lakens, D., & Leys, C. (2017). Why psychologists should by default use Welch’s *t*-test instead of Student’s *t*-test. *International Review of Social Psychology*, 30(1), 92–101.
- Drechsler, I., & Yaron, A. (2011). What’s vol got to do with it. *Review of Financial Studies*, 24(1), 1–45.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
- Engle, R. F., & Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *Journal of Finance*, 48(5), 1749–1778.
- Fleming, J., Ostdiek, B., & Whaley, R. E. (1995). Predicting stock market volatility: A new measure. *Journal of Futures Markets*, 15(3), 265–302.
- Fleming, J., Kirby, C., & Ostdiek, B. (2001). The economic value of volatility timing. *Journal of Finance*, 56(1), 329–352.
- Fousekis, P. (2020). Sign and size asymmetry in the stock returns–implied volatility relationship. *The Journal of Economic Asymmetries*, 21, e00162.

- Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48(5), 1779–1801.
- Haas, M., Mittnik, S., & Paolella, M. S. (2004). A new approach to Markov-switching GARCH models. *Journal of Financial Econometrics*, 2(4), 493–530.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.
- Harvey, C. R., Hoyle, E., Korgaonkar, R., Rattray, S., Sargaison, M., & Van Hemert, O. (2018). The impact of volatility targeting. *Journal of Portfolio Management*, 45(1), 14–33.
- Hibbert, A. M., Daigler, R. T., & Dupoyet, B. (2008). A behavioral explanation for the negative asymmetric return–volatility relation. *Journal of Banking & Finance*, 32(10), 2254–2266.
- Just, M., & Echaust, K. (2020). Stock market returns, volatility, correlation and liquidity during the COVID-19 crisis: Evidence from the Markov switching approach. *Finance Research Letters*, 37, 101775.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
- Lin, C.-C., Chen, C.-S., & Hwang, D.-Y. (2020). Does VIX or volume improve the prediction of intraday returns for the Taiwan stock market? *Pacific-Basin Finance Journal*, 61, 101316.
- Low, C. (2004). The fear and exuberance from implied volatility of S&P 100 index options. *Journal of Business*, 77(3), 527–546.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36(4), 394–419.

- Martin, I. (2017). What is the expected return on the market? *Quarterly Journal of Economics*, 132(1), 367–433.
- Moreira, A., & Muir, T. (2017). Volatility-managed portfolios. *Journal of Finance*, 72(4), 1611–1644.
- Muler, N., & Yohai, V. J. (2008). Robust estimates for GARCH models. *Journal of Statistical Planning and Inference*, 138(10), 2918–2940.
- Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2), 347–370.
- Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1), 246–256.
- Ruxton, G. D. (2006). The unequal variance t -test is an underused alternative to Student’s t -test and the Mann–Whitney U test. *Behavioral Ecology*, 17(4), 688–690.
- Todorov, V. (2010). Variance risk-premium dynamics: The role of jumps. *Review of Financial Studies*, 23(1), 345–383.
- Vlastakis, N., & Markellos, R. N. (2012). Information demand and stock market volatility. *Journal of Banking & Finance*, 36(6), 1808–1821.
- Whaley, R. E. (2000). The investor fear gauge. *Journal of Portfolio Management*, 26(3), 12–17.
- Zakoian, J.-M. (1994). Threshold heteroskedastic models. *Journal of Economic Dynamics and Control*, 18(5), 931–955.
- Zimmerman, D. W. (2004). A note on preliminary tests of equality of variances. *British Journal of Mathematical and Statistical Psychology*, 57(1), 173–181.

A Variable Definitions

Table 8: Variable Definitions

Variable	Definition
r_t	Daily log return: $\ln(P_t/P_{t-1})$
V_t	VIX closing level on day t
ΔV_t	Daily VIX change: $V_t - V_{t-1}$
NSI_t	Normalized shock impact: $ r_t /V_t$
SAR_j	Shock amplification ratio in regime j : $\overline{ r }_{S_j}/\overline{ r }_{N_j}$
VRP_t	Variance risk premium: $VIX_t^2/252 - RV_t^{(22)}/22$
$RV_t^{(h)}$	Realized variance over h days: $\sum_{i=1}^h r_{t-h+i}^2$

B Additional Cross-Asset Results

Table 9 reports the full absorption regression results for each asset, including intercepts and R^2 values.

Table 9: Full Absorption Regression Results by Asset

Asset	$\hat{\alpha}$	$\hat{\beta} (\times 10^4)$	$t(\hat{\beta})$	Adj R^2	N
SPY	0.0822	-2.73	-1.85	0.0076	769
GLD	0.0548	-4.34	-2.90	0.0142	769
TLT	0.0531	-4.37	-3.31	0.0290	769
0050.TW	0.0473	+0.92	+0.28	-0.0013	614

Notes: Dependent variable: $NSI_t = |r_t|/V_t$ with r_t in percent. Independent variable: V_t (VIX level). Newey-West standard errors with 10 lags. $\hat{\beta}$ is multiplied by 10^4 for readability. 0050.TW has fewer shock-day observations due to different trading calendar. Sample: 2006–2026 under the 2026-04-19 pinned data snapshot (`auto_adjust=False`); SPY/GLD/TLT/VIX from the pinned CSV, 0050.TW from yfinance with `auto_adjust=False`. SPY estimate ($\hat{\beta} \times 10^4 = -2.73$, $t = -1.85$, $N = 769$) is consistent with the SPY main-text pinned-snapshot footnote (-2.67 , $t = -1.77$, $N = 768$) up to a one-day sample-end boundary difference. Under the original paper-drafting snapshot (2025 yfinance pull), the corresponding estimates were SPY (-2.8 , $t = -3.42$), GLD (-4.3 , $t = -4.17$), TLT (-4.4 , $t = -3.89$), 0050.TW ($+1.9$, $t = +1.62$) with $N = 893/612$; signs are preserved across all four assets and U.S.-asset magnitudes are essentially unchanged, while t -statistics weaken due to the yfinance retroactive adjustment of VIX historical data (see Section 7 and Section 7 for the full snapshot-sensitivity disclosure). The primary SAR evidence (Table 3) and the RV-normalized regression ($\hat{\beta}^{\text{RV}} = -0.01249$, $t = -8.2$) are unaffected.