

Article

Not peer-reviewed version

---

# PCA-APT Stress Index for Market Drawdowns

---

[Ting Liu](#)\*

Posted Date: 5 March 2026

doi: 10.20944/preprints202603.0395.v1

Keywords: market stress; drawdown early warning; PCA; APT; sector ETFs; leakage-safe backtesting; risk management; volatility baseline



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# PCA-APT Stress Index for Market Drawdowns

Ting Liu

Northwestern University, Evanston, IL, USA; tingliu2027@u.northwestern.edu

## Abstract

This study develops a leakage-safe PCA–APT framework that constructs an idiosyncratic market-stress index from cross-sectional residual dispersion and evaluates its usefulness for anticipating equity drawdowns. Using daily adjusted prices for SPY and 11 U.S. sector ETFs from 2020–2025, we compute sector excess returns (sector minus SPY), estimate a low-dimensional common component via principal component analysis (PCA), and define residual stress as the cross-sectional root-mean-square magnitude of PCA reconstruction residuals. To prevent look-ahead bias, the PCA mapping is estimated using information available only through  $t-1$ , stress is computed out-of-sample at  $t$ , and stress regimes are identified using a rolling train-only quantile threshold that is shifted forward by one trading day. Drawdown-warning performance is assessed using drawdown-onset events and early-warning classification metrics (ROC-AUC, PR-AUC, and horizon- $H$  precision/recall). Empirically, residual stress spikes cluster around drawdown onsets and provides predictive information, although a volatility-based benchmark remains stronger on average across discrimination metrics. Importantly, residual stress exhibits *state-dependent complementarity* with volatility: conditional on *low* volatility, high residual stress is associated with a materially higher probability of a drawdown onset within the next  $H = 21$  trading days (approximately 17% vs. 8%), and the joint high-stress/high-volatility regime identifies the highest-risk states (approximately 36% onset probability). Event-level overlap diagnostics further indicate that residual stress can flag a subset of drawdown onsets not captured by a volatility-threshold rule, while some onsets are not preceded by either signal. Economic relevance is examined under transaction costs through (i) a residual-ranked sector long–short portfolio and (ii) stress-managed SPY overlays that reduce exposure during detected regimes. In the baseline sample, a volatility-managed overlay improves drawdown control relative to buy-and-hold, whereas the residual-stress overlay does not reduce maximum drawdown and the residual-ranked long–short strategy is not robustly profitable after costs. Overall, the paper contributes a reproducible, leakage-safe evaluation pipeline linking cross-sectional residual dispersion to drawdown risk and clarifies when residual stress serves as a complementary market-structure risk indicator alongside standard volatility-based signals.

**Keywords:** market stress; drawdown early warning; PCA; APT; sector ETFs; leakage-safe backtesting; risk management; volatility baseline

## 1. Introduction

Large equity drawdowns are economically costly and difficult to forecast with stable linear predictors. A practical alternative is *state detection*: identifying regimes in which downside risk is elevated so that investors can adapt exposure and risk limits. This paper develops a simple, interpretable, and leakage-safe market-stress indicator derived from the cross-section of U.S. sector ETF returns and evaluates its usefulness for anticipating drawdown onsets.

Our starting point is that sector returns often admit a low-dimensional common structure driven by macroeconomic and market-wide forces. During transitions into risk-off conditions, sector repricing and rotation can become less synchronized, increasing *idiosyncratic* cross-sectional dispersion beyond what is explained by dominant common components [4–6]. We capture this mechanism by applying

principal component analysis (PCA) [3] to *sector excess returns* (sector minus SPY), extracting a  $K$ -dimensional common component, and measuring stress as the cross-sectional root-mean-square magnitude of *PCA reconstruction residuals*. Intuitively, the stress score rises when sector moves are increasingly “out of line” with the low-dimensional market-driven structure [2].

A core contribution is methodological: we emphasize *leakage-safe* regime construction. At each date  $t$ , the PCA mapping (centering and loadings) is estimated using information available only through  $t-1$ , and the residual stress score is computed out-of-sample for date  $t$ . Stress regimes are then labeled using a rolling *train-only* quantile threshold computed from historical stress scores and shifted forward by one trading day, so that the regime label used for trading at  $t$  does not incorporate information from  $t$  itself. Trading signals are lagged accordingly, enabling a clean evaluation of both early-warning performance and net-of-cost strategy behavior.

This study evaluates drawdown-warning utility using drawdown-onset events and classification-style metrics (ROC-AUC, PR-AUC, and horizon- $H$  precision/recall), complemented by event-study visualizations around onset dates. This study further assesses economic relevance through two applications with transaction costs: (i) a residual-ranked sector long-short portfolio and (ii) a stress-managed SPY overlay that reduces exposure during detected stress regimes. Because volatility-based measures are a natural benchmark for market stress [7,8], this study also compares the residual-based stress signal to a volatility baseline within the same leakage-safe timing protocol.

The paper makes four contributions:

- **Residual stress index:** We construct a leakage-safe, PCA-based residual stress score defined as the cross-sectional root-mean-square magnitude of out-of-sample PCA reconstruction residuals from sector excess returns (sector minus SPY), capturing idiosyncratic dispersion beyond the top- $K$  common components.
- **Leakage-safe regime labeling:** We propose an implementable, train-only quantile thresholding scheme in which (i) PCA estimation and stress scoring use information available only through  $t-1$ , and (ii) the rolling threshold is shifted forward by one trading day so that the regime label at  $t$  depends only on information available through  $t-1$ .
- **Drawdown-onset early-warning evaluation:** We evaluate warning utility using drawdown-onset events and horizon- $H$  early-warning metrics (ROC-AUC, PR-AUC, and precision/recall), complemented by event-study diagnostics and regime-conditional onset rates.
- **Economic relevance, benchmarking, and complementarity:** This study connects the signal to practice via two transaction-cost-aware applications—(i) a residual-ranked sector long-short portfolio and (ii) a stress-managed SPY overlay for downside risk control—and benchmark residual stress against a volatility-based baseline under the same leakage-safe timing protocol. This study further quantifies when residual stress provides incremental, state-dependent information beyond volatility using joint-regime (stress/volatility) diagnostics.

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 describes the data, leakage-safe signal construction, and drawdown-onset definition. Section 4 reports early-warning performance, robustness checks, and trading applications. Section 5 discusses interpretation and limitations. Section 6 concludes.

Volatility is a strong benchmark; residual stress is evaluated as a complementary lens and an interpretability tool, not guaranteed to dominate.

## 2. Literature Review

This paper relates to three strands of research: (i) factor models and dimensionality reduction for asset returns, (ii) stress and tail-risk measurement using cross-sectional information, and (iii) early-warning evaluation and risk-managed overlays for drawdown control.

### 2.1. Factor Structure, PCA, and APT-Style Decompositions

A large body of work models asset returns using a small number of common factors. The Arbitrage Pricing Theory (APT) provides a conceptual foundation for multi-factor representations of expected returns, while empirical asset-pricing research has developed widely used factor benchmarks. In parallel, statistical factor methods and dynamic factor models estimate latent factor spaces directly from large return panels when the number or identity of true factors is unknown. Within this literature, PCA is a common tool for extracting a low-rank representation of co-movement and isolating an orthogonal residual component [3,15]. In this study, PCA plays a dual role: it (i) extracts the dominant common component in sector *excess* returns (sector minus SPY), thereby removing the market-wide mode, and (ii) yields a clean residual panel that can be used to measure idiosyncratic dispersion beyond the top- $K$  common components [15,16].

### 2.2. Cross-Sectional Residual Dispersion as a Stress Signal

Beyond average co-movement, cross-sectional return *dispersion* and residual variation can carry information about market states. Intuitively, elevated residual dispersion reflects disagreement, dislocation, or sector rotation that is not captured by broad market variation [4–6]. Related ideas appear in work on cross-sectional dispersion, disagreement, and uncertainty measures constructed from panels of asset returns. Our approach operationalizes this concept through a transparent residual stress score: the cross-sectional root-mean-square magnitude of PCA reconstruction residuals from sector excess returns. Because volatility-based measures (realized volatility, implied volatility, and related proxies) are widely used as market stress indicators [7,8], we treat them as natural benchmarks for evaluating whether residual dispersion provides complementary early-warning information.

### 2.3. Drawdown Risk, Early-Warning Evaluation, and Risk-Managed Overlays

Predicting drawdowns is challenging because downside tail events are rare and market dynamics are often state-dependent [11,18]. A common response is to frame the problem as *regime detection* or *early warning*: signals are assessed using event-based definitions and classification-style metrics rather than only mean-return predictability. In portfolio applications, regime signals are frequently mapped into overlays that reduce risk exposure during stress periods, trading off protection against tracking error, turnover, and transaction costs [11]. A related line of work studies volatility-managed or risk-controlled exposures as implementable mechanisms for drawdown mitigation [7,8]. Consistent with this perspective, we treat residual stress primarily as a *regime indicator* and evaluate both its warning utility (drawdown-onset early warning using ROC-AUC, PR-AUC, and horizon- $H$  precision/recall) and its economic relevance (net-of-cost overlay behavior under lagged, implementable timing) [9,10,18].

### 2.4. Positioning of This Work

Relative to prior work, our contribution is best understood as an evaluation framework that is both economically interpretable and explicitly leakage-safe. The distinguishing features are threefold. First, we construct an idiosyncratic stress score from sector *excess* returns (sector minus SPY), removing the contemporaneous market component and quantifying cross-sectional residual dispersion beyond a low-rank common factor space. Second, we enforce a *leakage-safe* timing protocol throughout: the PCA mapping and residual stress at date  $t$  are estimated using information available only through  $t-1$ , stress regimes are identified via a rolling *train-only* quantile threshold that is shifted forward by one trading day, and any trading rule uses lagged signals. Third, we provide a unified empirical assessment that combines drawdown-onset early-warning metrics and event-study diagnostics with transaction-cost-aware economic tests, while benchmarking residual stress against a realized-volatility baseline under the same leakage-safe protocol. This design clarifies when residual stress complements—rather than necessarily dominates—standard volatility-based risk signals in practical drawdown monitoring.

### 3. Methodology

#### 3.1. Data Preparation and Processing

This study uses daily adjusted close prices for the S&P 500 ETF (SPY) and 11 U.S. sector ETFs: XLC, XLY, XLP, XLE, XLF, XLV, XLI, XLK, XLB, XLU, and XLRE. All series are retrieved from Yahoo Finance over the sample period 2020-01-01 to 2025-12-31 (the final observation date in our downloaded snapshot). Adjusted prices are used to account for dividends and stock splits. After download, all series are aligned to a common trading calendar; any date with missing observations in any ETF is removed to ensure a balanced panel. Because Yahoo Finance data may be updated over time, the results reported in this paper correspond to the specific snapshot used in our empirical analysis, which ends on 2025-12-31; the replication script and configuration are provided to reproduce this endpoint.

#### 3.2. Data and Code Availability

All data used in this study are publicly available. Daily adjusted close prices for SPY and the 11 U.S. sector ETFs are retrieved from Yahoo Finance via quantmod. A complete replication package (R scripts and configuration files) that reproduces all tables and figures is provided as Supplementary Materials accompanying this article. The configuration file freezes the sample period to 2020-01-01 through 2025-12-31 and preserves the leakage-safe timing protocol used throughout the paper. The package does not redistribute the raw price data; instead, the download script retrieves them directly from the public source.

#### 3.3. Returns and Sector Excess Returns

Let  $P_{i,t}$  denote the adjusted close price of sector ETF  $i$  on day  $t$ , and let  $P_{m,t}$  denote the adjusted close price of the market ETF (SPY). We compute simple daily returns as

$$r_{i,t} = \frac{P_{i,t}}{P_{i,t-1}} - 1, \quad r_{m,t} = \frac{P_{m,t}}{P_{m,t-1}} - 1. \quad (1)$$

To isolate sector-specific movements beyond the broad market component, we define sector excess returns (sector minus SPY) as

$$x_{i,t} = r_{i,t} - r_{m,t}. \quad (2)$$

Stacking  $x_{i,t}$  across the  $N = 11$  sectors yields the excess-return vector  $\mathbf{x}_t \in \mathbb{R}^N$ . Collecting observations over  $t = 1, \dots, T$  forms the return panel

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1^\top \\ \vdots \\ \mathbf{x}_T^\top \end{bmatrix} \in \mathbb{R}^{T \times N}, \quad (3)$$

which serves as the input to the PCA-APT decomposition in the next subsection.

#### 3.4. PCA Factor Extraction and Residual Stress Index

Let  $x_{i,t}$  denote the sector excess return (sector  $i$  minus SPY) on day  $t$  and let  $\mathbf{x}_t = (x_{1,t}, \dots, x_{N,t})^\top \in \mathbb{R}^N$ . Collecting observations over  $t = 1, \dots, T$  forms the return panel  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]^\top \in \mathbb{R}^{T \times N}$ . PCA provides a low-rank representation of cross-sectional co-movement by projecting  $\mathbf{X}$  onto the top- $K$  principal components. Denote by  $\mathbf{L} \in \mathbb{R}^{N \times K}$  the loading matrix and by  $\mathbf{F} \in \mathbb{R}^{T \times K}$  the factor-score matrix. The rank- $K$  reconstruction of  $\mathbf{X}$  is

$$\widehat{\mathbf{X}} = \mathbf{F}\mathbf{L}^\top, \quad (4)$$

and the corresponding residual matrix is

$$\mathbf{E} = \mathbf{X} - \widehat{\mathbf{X}}, \quad (5)$$

where  $e_{i,t}$  captures the sector-specific deviation from the low-dimensional common component at time  $t$ . This study defines the PCA–APT residual stress score as the cross-sectional root-mean-square (RMS) magnitude of these residuals,

$$\text{Stress}_t = \sqrt{\frac{1}{N} \sum_{i=1}^N e_{i,t}^2}. \quad (6)$$

Intuitively,  $\text{Stress}_t$  increases when sector returns become less synchronized and many sectors deviate from the dominant common structure implied by the top- $K$  components.

**Leakage-safe out-of-sample stress scoring.** To avoid look-ahead bias, both the PCA mapping and the stress score are constructed out-of-sample. At each date  $t$ , the PCA loadings and centering parameters are estimated using only information available through  $t-1$  (either an expanding window or a rolling window of length  $W_{\text{PCA}}$ ), and the residual vector  $\mathbf{e}_t = (e_{1,t}, \dots, e_{N,t})^\top$  is computed by projecting  $\mathbf{x}_t$  onto the top- $K$  loadings obtained from the training sample and taking the reconstruction error. The stress score  $\text{Stress}_t$  in (6) is then computed from  $\mathbf{e}_t$ . This timing ensures that  $\text{Stress}_t$  is implementable at time  $t$  without using information from future dates.

**Out-of-sample regime labeling (train-only threshold).** Stress regimes are labeled using a rolling train-only quantile threshold computed from historical stress scores and shifted forward by one trading day.

### 3.5. PCA–APT Projection and Residual Construction

Let  $\mathbf{x}_t \in \mathbb{R}^N$  denote the vector of sector excess returns on day  $t$  and let  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]^\top \in \mathbb{R}^{T \times N}$ . We center the panel using training-sample moments. Specifically, let  $\boldsymbol{\mu}_{t-1} \in \mathbb{R}^N$  denote the sample mean of  $\mathbf{x}_s$  computed using only  $\{s \leq t-1\}$  (either expanding or rolling, consistent with the PCA estimation window). For the baseline specification reported in the main tables and figures, we use an *expanding-window* PCA: at each date  $t$  the factor space is estimated using all sector excess returns available up to  $t-1$  (starting from an initial burn-in window of  $W_{\text{PCA}} = 252$  trading days). We define the centered vector

$$\tilde{\mathbf{x}}_t = \mathbf{x}_t - \boldsymbol{\mu}_{t-1}. \quad (7)$$

**PCA factor space (estimated with information through  $t-1$ ).** Let  $\mathbf{L}_{t-1} = [\boldsymbol{\ell}_{1,t-1}, \dots, \boldsymbol{\ell}_{K,t-1}] \in \mathbb{R}^{N \times K}$  denote the orthonormal loading matrix obtained by applying PCA to the centered training panel  $\{\mathbf{x}_s - \boldsymbol{\mu}_{t-1}\}_{s \leq t-1}$ , so that  $\mathbf{L}_{t-1}^\top \mathbf{L}_{t-1} = \mathbf{I}_K$ . The corresponding (out-of-sample) factor score for day  $t$  is

$$\mathbf{f}_t = \mathbf{L}_{t-1}^\top \tilde{\mathbf{x}}_t \in \mathbb{R}^K. \quad (8)$$

**APT projection (rank- $K$  reconstruction) and residuals.** The PCA–APT common component is the orthogonal projection of  $\tilde{\mathbf{x}}_t$  onto the span of the first  $K$  principal components estimated from the training sample:

$$\hat{\mathbf{x}}_t = \mathbf{L}_{t-1} \mathbf{f}_t = \mathbf{L}_{t-1} \mathbf{L}_{t-1}^\top \tilde{\mathbf{x}}_t. \quad (9)$$

This study defines the residual vector as the reconstruction error,

$$\mathbf{e}_t = \tilde{\mathbf{x}}_t - \hat{\mathbf{x}}_t = (\mathbf{I}_N - \mathbf{L}_{t-1} \mathbf{L}_{t-1}^\top) \tilde{\mathbf{x}}_t \in \mathbb{R}^N. \quad (10)$$

By construction,  $\mathbf{e}_t$  is orthogonal to the estimated factor space and isolates sector-specific deviations not explained by the top- $K$  common components.

$$\mathbf{L}_{t-1}^\top \mathbf{e}_t = \mathbf{0}, \quad (11)$$

**Least-squares optimality (training sample).** For a fixed training window, the rank- $K$  reconstruction obtained by PCA minimizes the Frobenius loss on the centered training panel, i.e., it solves  $\min_{\text{rank}(\mathbf{Y}) \leq K} \|\tilde{\mathbf{X}}_{\text{train}} - \mathbf{Y}\|_F^2$ . This standard property justifies interpreting  $\hat{\mathbf{x}}_t$  as the best low-rank approximation of common co-movement, and  $\mathbf{e}_t$  as the corresponding idiosyncratic component.

### 3.6. Residual Stress Score

The out-of-sample residual vector  $\mathbf{e}_t$  is defined in Eq. (10) and is orthogonal to the estimated factor space (Eq. (11)). We summarize the cross-sectional magnitude of these idiosyncratic deviations using the RMS residual stress score defined in Eq. (6). Intuitively,  $\text{Stress}_t$  increases when sector moves become less synchronized and the low-rank common structure provides a poorer reconstruction of the cross section.

### 3.7. Walk-Forward, Leakage-Safe Construction, and Implementable Signals

To ensure a leakage-safe (out-of-sample) design, we align all price series to a common trading-day calendar, compute returns, and form the sector excess-return panel  $\mathbf{X} \in \mathbb{R}^{T \times N}$  as in Section 3.3. Throughout, any quantity used for labeling or trading at date  $t$  is constructed using information available no later than  $t - 1$ .

**Train-only rolling threshold (shifted by one day).** Let  $\text{Stress}_t$  denote the leakage-safe residual stress score defined in Eq. (6).

Using a lookback window of length  $W_{\text{thr}}$  and quantile level  $q$ , we compute the train-only threshold

$$Q_{t-1}^{(q)} = \text{Quantile}_q(\text{Stress}_{t-W_{\text{thr}}}, \dots, \text{Stress}_{t-1}), \quad (12)$$

and define the stress-regime indicator as

$$\mathbb{I}_t^{\text{stress}} = \mathbb{1}\{\text{Stress}_t > Q_{t-1}^{(q)}\}. \quad (13)$$

In implementation, the rolling quantile in Eq. (12) is evaluated with a right-aligned window and then shifted forward by one trading day, so that the threshold applied at time  $t$  depends only on  $\{\text{Stress}_\tau\}_{\tau \leq t-1}$ .

**Lagged trading signals (implementable timing).** Any strategy component that depends on the stress regime (e.g., the SPY overlay) uses the lagged regime label  $\mathbb{I}_{t-1}^{\text{stress}}$  to set exposure for day  $t$  (one-day signal lag). For the residual-ranked long-short strategy, sector ranking signals are formed from the lagged residual vector  $\mathbf{e}_{t-1}$ , portfolio weights are set at the close of  $t - 1$ , and realized portfolio returns are computed using contemporaneous excess returns  $\mathbf{x}_t$ . This timing ensures that both regime identification and portfolio formation are implementable and free of look-ahead bias.

### 3.8. Drawdown-Onset Events and Horizon- $H$ Early-Warning Labels

**SPY drawdown and drawdown-onset events.** Let  $r_{m,t}$  denote the simple daily return of SPY (Eq. (1)) and define the cumulative equity curve  $E_t = \prod_{\tau \leq t} (1 + r_{m,\tau})$ . The drawdown series is computed as

$$\text{DD}_t = \frac{E_t}{\max_{\tau \leq t} E_\tau} - 1. \quad (14)$$

Fix a drawdown threshold  $\delta < 0$  (e.g.,  $\delta = -10\%$ ). We define the drawdown-region indicator as

$$\mathbb{I}_t^{\text{DD}} = \mathbb{1}\{\text{DD}_t \leq \delta\}, \quad (15)$$

and define a drawdown-onset event as the first entry into the drawdown region:

$$\mathbb{I}_t^{\text{onset}} = \mathbb{I}_t^{\text{DD}} (1 - \mathbb{I}_{t-1}^{\text{DD}}). \quad (16)$$

Thus,  $\mathbb{I}_t^{\text{onset}} = 1$  only on dates when the drawdown crosses below  $\delta$  from above, avoiding repeated labels on consecutive days within the same drawdown episode.

**Horizon- $H$  early-warning label.** For classification-style early warning, we define the binary target

$$y_t^{(H)} = \mathbb{1}\left\{\sum_{h=1}^H \mathbb{1}_{t+h}^{\text{onset}} \geq 1\right\}, \quad (17)$$

which equals one if at least one drawdown-onset event occurs within the next  $H$  trading days. The unconditional event rate used in precision–recall analysis is

$$\bar{y} = \frac{1}{T} \sum_{t=1}^T y_t^{(H)} = \Pr(y_t^{(H)} = 1). \quad (18)$$

We pre-specify sensitivity checks over  $K$ ,  $q$ ,  $W_{\text{thr}}$ ,  $\delta$ , and  $H$ ; all runs preserve the same leakage-safe timing protocol (Section 3.7).

### 3.9. Realized-Volatility Baseline (SPY)

**Volatility score.** Let  $r_{m,t}$  denote the simple daily return of SPY (Eq. (1)). We define realized volatility as the annualized rolling standard deviation of daily returns over a window of length  $W_{\text{vol}}$ :

$$\text{Vol}_t = \sqrt{252} \cdot \text{SD}(r_{m,t-W_{\text{vol}}+1}, \dots, r_{m,t}). \quad (19)$$

The continuous score  $\text{Vol}_t$  is used as the volatility benchmark in ROC/PR evaluation.

**Volatility regime label (train-only threshold).** Analogous to the stress-regime construction, we compute a train-only rolling quantile threshold using a lookback window of length  $W_{\text{thr,vol}}$  and quantile level  $q_{\text{vol}}$ :

$$Q_{t-1, \text{vol}}^{(q_{\text{vol}})} = \text{Quantile}_{q_{\text{vol}}}(\text{Vol}_{t-W_{\text{thr,vol}}}, \dots, \text{Vol}_{t-1}), \quad (20)$$

and define the volatility-regime indicator as

$$\mathbb{I}_t^{\text{vol}} = \mathbb{1}\left\{\text{Vol}_t > Q_{t-1, \text{vol}}^{(q_{\text{vol}})}\right\}. \quad (21)$$

In implementation, the rolling quantile in Eq. (20) is evaluated with a right-aligned window and shifted forward by one trading day so that the regime label at date  $t$  depends only on information available through  $t - 1$ .

### 3.10. Volatility/Stress Overlay and Transaction Costs

**Baseline strategy return.** Let  $\mathbf{w}_{t-1}$  denote the portfolio weights set at the close of day  $t - 1$  (held over day  $t$ ), and let  $\mathbf{x}_t$  denote the vector of contemporaneous simple excess returns of the traded assets on day  $t$ . The gross (pre-cost) portfolio return is

$$r_{p,t}^{\text{gross}} = \mathbf{w}_{t-1}^\top \mathbf{x}_t. \quad (22)$$

**Overlay rule (“30% overlay”).** Let  $\mathbb{I}_{t-1}^{\text{ov}} \in \{0, 1\}$  be an implementable overlay trigger known at time  $t - 1$  (e.g.,  $\mathbb{I}_{t-1}^{\text{stress}}$  or  $\mathbb{I}_{t-1}^{\text{vol}}$ ). For an overlay intensity  $\lambda \in (0, 1)$  (e.g.,  $\lambda = 0.30$ ), we scale the risky exposure by a factor

$$s_{t-1} = 1 - \lambda \mathbb{I}_{t-1}^{\text{ov}}. \quad (23)$$

The overlay-adjusted weights are  $\tilde{\mathbf{w}}_{t-1} = s_{t-1} \mathbf{w}_{t-1}$ , so a “30% overlay” reduces risky exposure to 70% of the baseline weights when the trigger is active and leaves the portfolio unchanged otherwise. The remaining fraction  $1 - s_{t-1}$  is allocated to cash with zero return, so the overlay-adjusted gross return is

$$\tilde{r}_{p,t}^{\text{gross}} = \tilde{\mathbf{w}}_{t-1}^\top \mathbf{x}_t = s_{t-1} \mathbf{w}_{t-1}^\top \mathbf{x}_t. \quad (24)$$

**Transaction costs and net returns.** We apply proportional transaction costs to portfolio turnover. Let

$$\text{TO}_t = \sum_i |\tilde{w}_{i,t} - \tilde{w}_{i,t-1}| \quad (25)$$

denote one-way turnover in weights at the rebalance from  $t - 1$  to  $t$  (computed after applying the overlay). Given a one-way cost rate  $c$  (in decimal form), the net return is

$$r_{p,t}^{\text{net}} = \tilde{r}_{p,t}^{\text{gross}} - c \cdot \text{TO}_t. \quad (26)$$

Unless otherwise stated, we set  $c = 5$  bps per unit of one-way turnover and report performance statistics using the net return series  $\{r_{p,t}^{\text{net}}\}$ .

## 4. Results

### 4.1. Descriptive PCA Structure in SECTOR Excess Returns

Figure 1 reports the PCA scree plot for the sector excess-return panel (sector minus SPY). Consistent with Table 1, the first two principal components explain the majority of the cross-sectional variation (about 63.8% cumulatively), while subsequent components contribute incrementally. This pattern supports the parsimonious choice  $K = 2$ , which we use for the subsequent PCA–APT projection and residual-based stress construction.

**Fig 1: PCA Scree (Sector Excess Returns)**


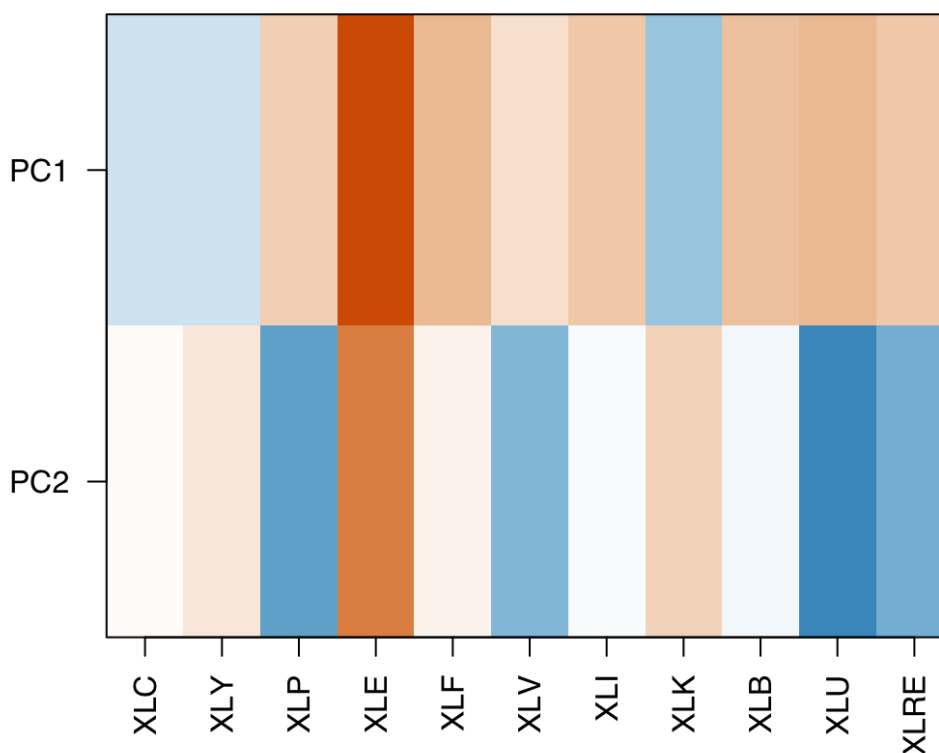


Figure 1. PCA scree plot of sector excess returns (sector minus SPY).

**Table 1.** Variance explained by the top principal components (descriptive PCA on sector excess returns).

PC	Explained variance (%)	Cumulative (%)
PC1	38.210	38.210
PC2	25.567	63.776
PC3	7.902	71.679
PC4	6.347	78.025
PC5	5.439	83.465
PC6	4.158	87.623
PC7	3.993	91.616
PC8	3.445	95.060
PC9	2.504	97.564
PC10	1.741	99.305


To interpret the economic content of the leading components, Figure 2 displays the PC1–PC2 loading heatmap across the 11 sector ETFs, and Table 2 reports the largest contributors by absolute loading. The heatmap indicates that a low-dimensional factor structure captures much of the common variation in sector excess returns. Differences in sign reflect opposite directional co-movement with the corresponding latent component.

**Fig 2: PCA Loadings Heatmap (Sectors)****Figure 2.** PC1–PC2 loading heatmap across the 11 sector ETFs.

**Table 2.** Largest absolute PCA loadings for PC1–PC2 (descriptive PCA).

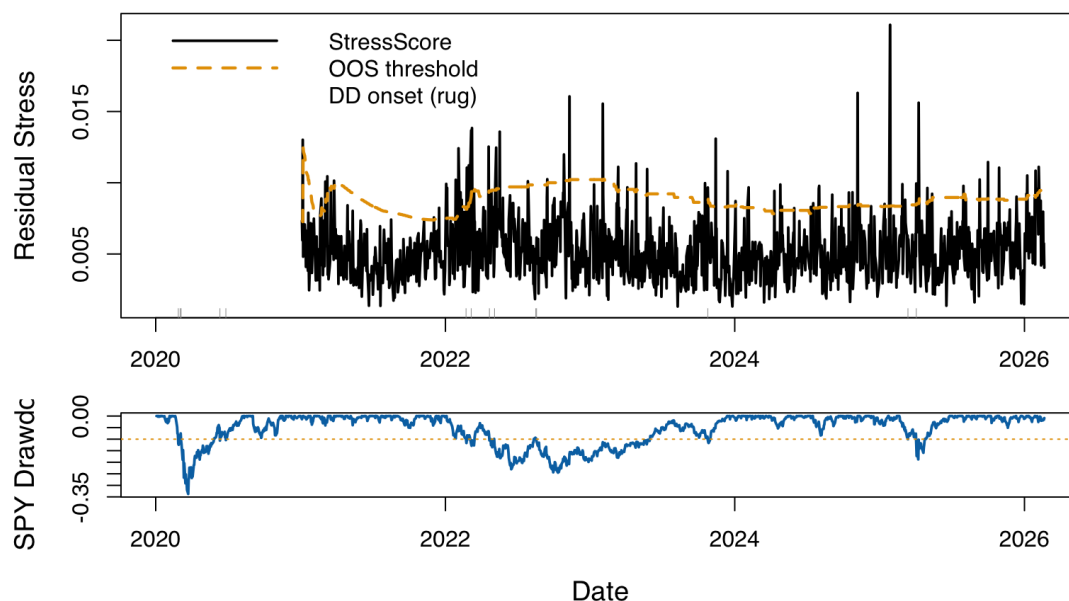
PC	Ticker	Loading
PC1	XLE	0.7544
PC1	XLF	0.2696
PC1	XLU	0.2673
PC1	XLK	-0.2580
PC2	XLU	-0.5379
PC2	XLE	0.5149
PC2	XLP	-0.4209
PC2	XLRE	-0.3625

Figure 3 plots the time series of the first two PCA factor scores ( $K = 2$ ). Shaded regions indicate dates classified as stress regimes using the leakage-safe, out-of-sample rolling quantile threshold applied to the residual stress score. These episodes correspond to intervals in which sector excess returns are less well explained by the low-dimensional common structure (i.e., residual dispersion is elevated) relative to its recent trailing distribution.

**Fig 3: PCA Factor Scores (K=2) with Stress Shading****Figure 3.** PC1–PC2 factor scores with stress periods shaded.

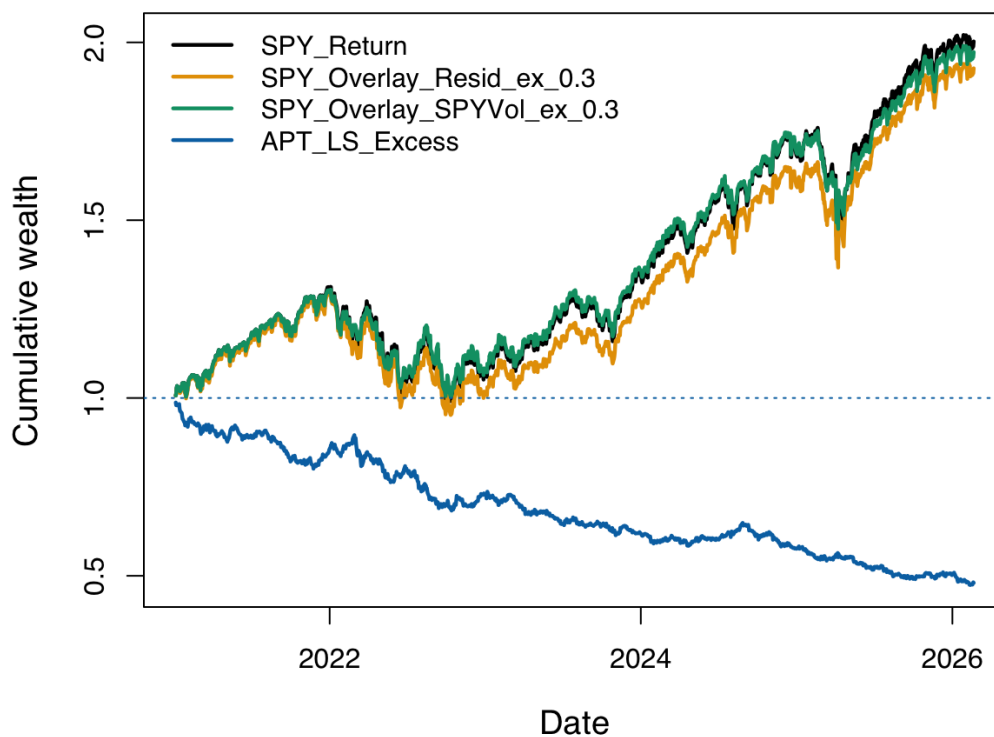
#### 4.2. Residual Stress Dynamics and Drawdowns

Figure 4 overlays the residual stress score with its train-only rolling quantile threshold (top panel) and compares it with the SPY drawdown series (bottom panel). Rug marks indicate drawdown-onset events (first entry below the drawdown threshold). Stress spikes frequently occur near these onset dates, indicating that elevated cross-sectional residual dispersion tends to coincide with—and in some episodes precede—the start of market drawdowns.

**Fig 4: Residual Stress vs SPY Drawdown****Figure 4.** Residual stress (top) and SPY drawdown (bottom) with out-of-sample thresholds.

#### 4.3. Economic Performance and Implementation Frictions

Figure 5 reports cumulative wealth for SPY buy-and-hold returns, two stress-managed SPY overlay variants (residual-stress and volatility-based), and the residual-ranked APT long-short portfolio (reported as a net return series). Figure 6 shows the corresponding drawdown profiles.

**Fig 5: Cumulative Wealth (Strategies)****Figure 5.** Cumulative wealth for SPY buy-and-hold, two stress-managed SPY overlays, and the APT residual-ranked long-short strategy (net returns).


**Fig 6: Drawdowns (Strategies)****Figure 6.** Drawdowns for SPY buy-and-hold and the stress-managed SPY overlay strategies (net returns).

Table 3 summarizes annualized performance over the sample. The volatility-based overlay attains the highest Sharpe ratio (0.944) and the smallest maximum drawdown (-0.231) among the SPY-based strategies. The residual-stress overlay slightly underperforms the SPY buy-and-hold baseline and exhibits a somewhat deeper drawdown. The residual-ranked APT long-short strategy performs poorly in this sample, consistent with residual signals being weak and sensitive to implementation costs.

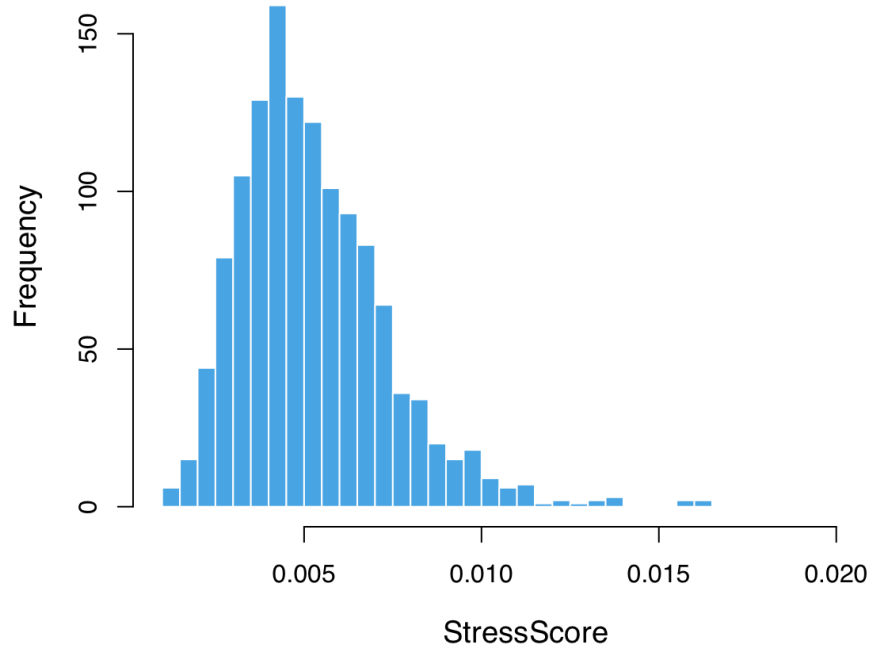

**Table 3.** Strategy performance summary (annualized).

Strategy	Ann. Return	Ann. Vol	Sharpe	MaxDD
SPY buy-and-hold	0.1455	0.1699	0.8848	-0.2450
SPY overlay (residual stress, 30%)	0.1369	0.1652	0.8592	-0.2658
SPY overlay (volatility, 30%)	0.1422	0.1532	0.9444	-0.2308
APT residual-ranked long-short	-0.1337	0.0998	-1.3878	-0.5198

Note: "Overlay (30%)" denotes scaling risky portfolio weights by 0.7 when the specified trigger is active (the remaining 30% is held in cash at zero return). "Net returns" subtract proportional transaction costs of  $c = 5$  bps per unit of one-way turnover.

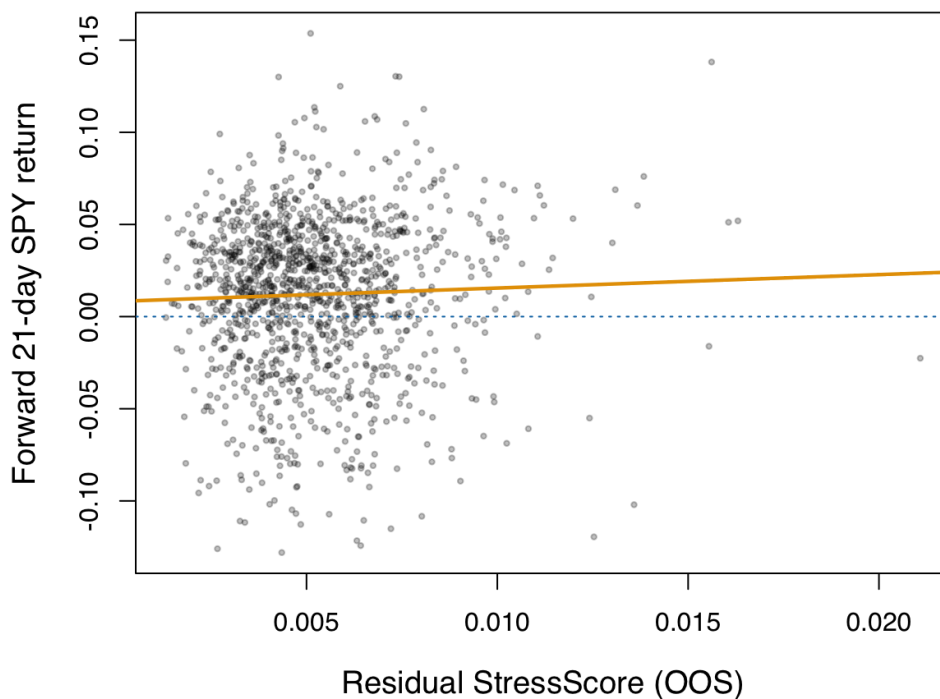
#### 4.4. Distributional Properties of Residual Stress

Figure 7 shows that the residual stress score is positively skewed with a pronounced right tail. The Q-Q plot in Figure 8 indicates departures from normality that are concentrated in the tails, consistent with occasional episodes of elevated cross-sectional residual dispersion.

**Fig 7: Residual Stress Distribution****Figure 7.** Distribution of the residual stress score.**Fig 8: Q-Q Plot (Residual Stress)****Figure 8.** Normal Q-Q plot of the residual stress score.

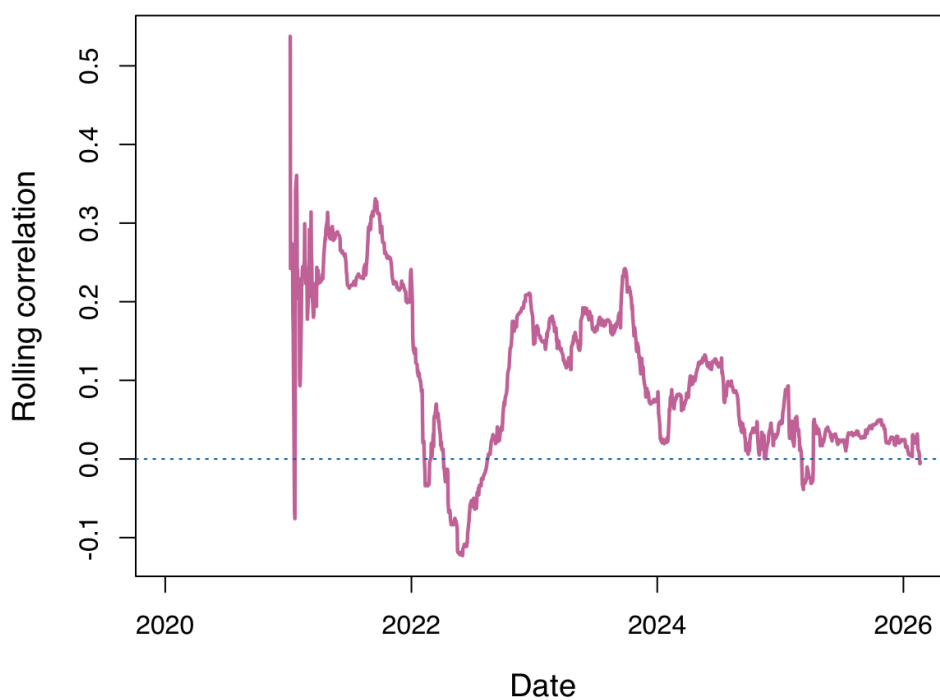
#### 4.5. Residual Stress and Forward SPY Returns

Figure 9 relates the residual stress score to the forward 21-day SPY return. The fitted line summarizes the average linear association, which appears weak in magnitude. This suggests that any predictive relationship is limited on average (and may be nonlinear or state-dependent).

**Fig 9: Residual Stress vs Forward 21d SPY Return**

**Figure 9.** Residual stress score (out-of-sample) versus forward 21-day SPY return; the line shows the fitted linear trend.

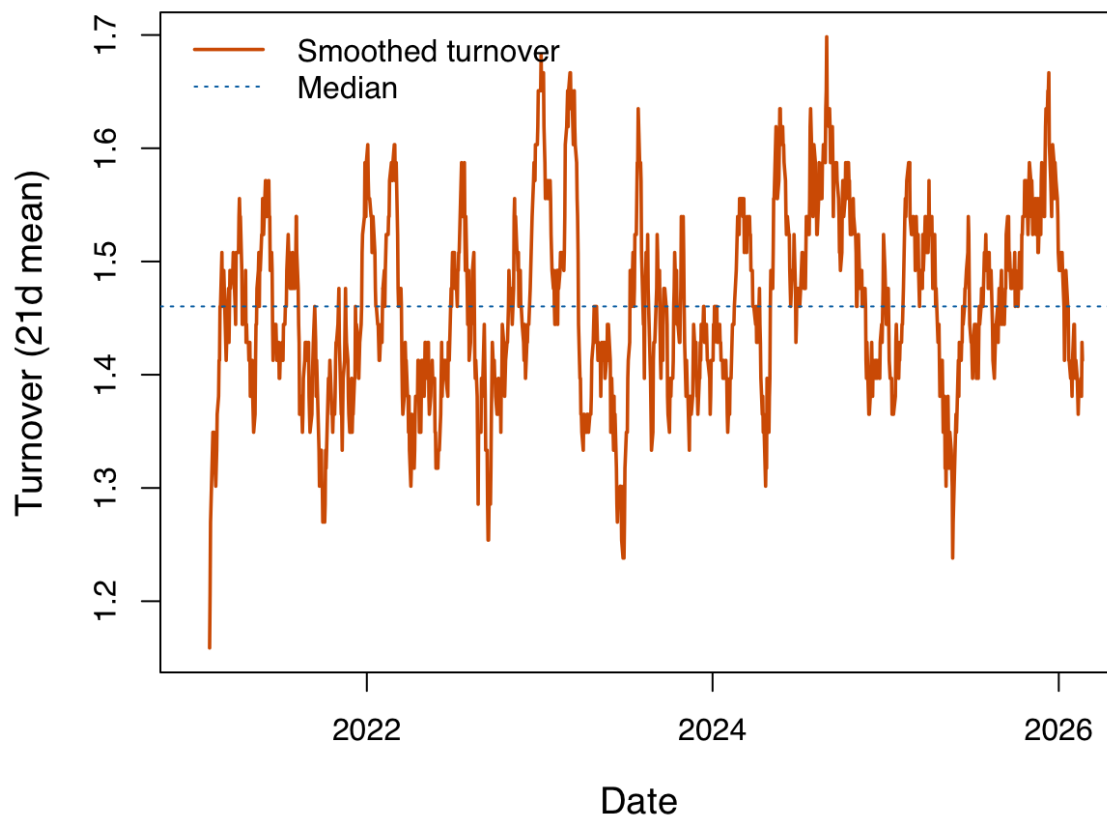
To examine time variation, Figure 10 reports the rolling 252-day correlation between the residual stress score and the forward 21-day SPY return. The correlation exhibits substantial time variation, suggesting that any linear association is episodic rather than persistent.

**Fig 10: Rolling Corr (252d) Stress vs Forward 21d Retu**

**Figure 10.** Rolling 252-day correlation between residual stress and forward 21-day SPY return.

Figure 11 reports the 21-day rolling mean of the APT long–short portfolio’s one-way turnover, defined as  $\frac{1}{2} \sum_i |w_{t,i} - w_{t-1,i}|$ . The persistently high turnover implies frequent rebalancing, so transaction costs can materially reduce net performance.

**Fig 11: APT LS Turnover (21d rolling mean)**

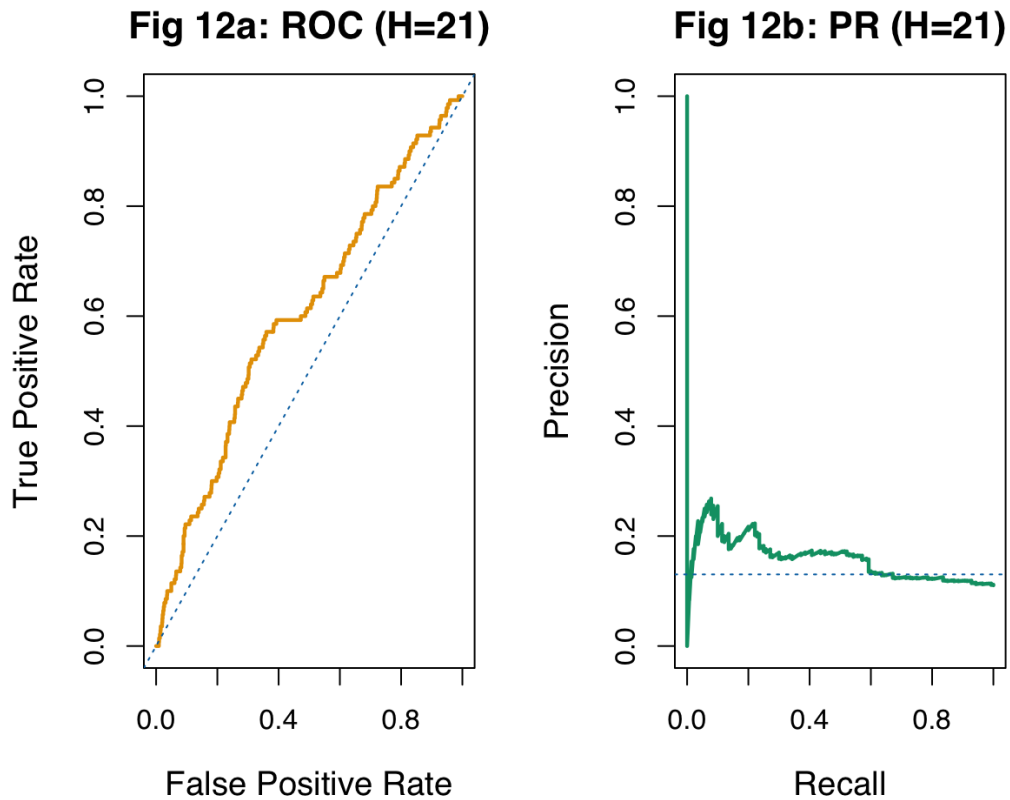


**Figure 11.** APT long–short turnover (21-day rolling mean).

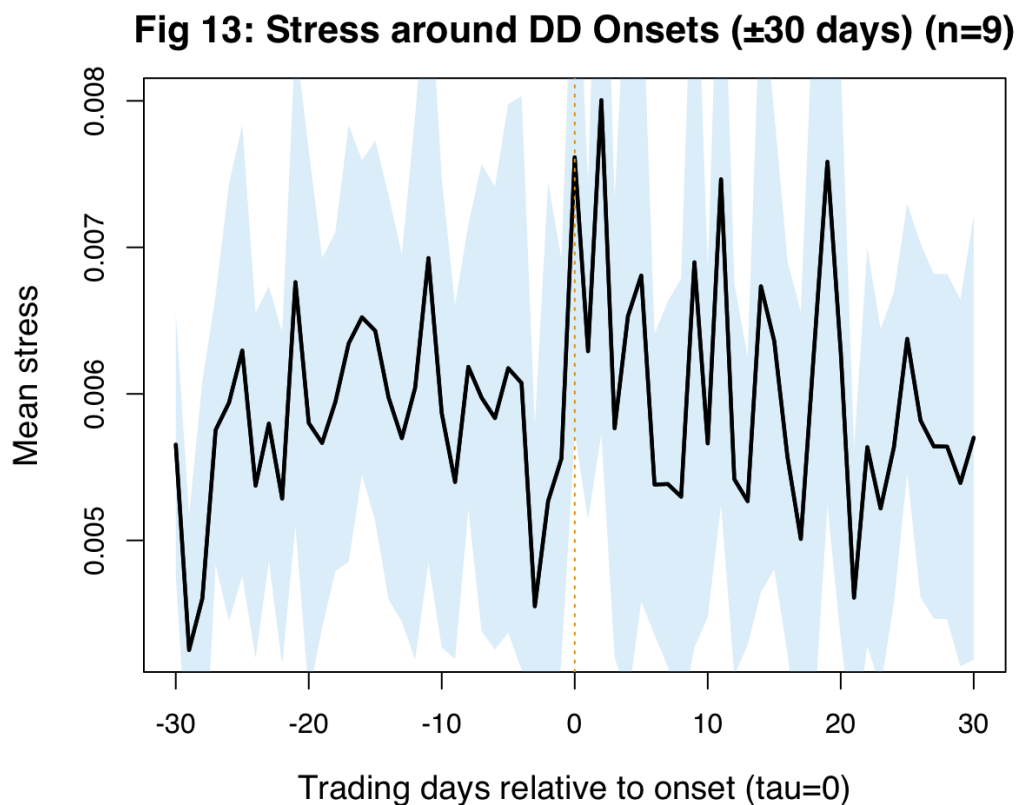
#### 4.6. Early-Warning Evaluation for Drawdown Onsets

Figure 12 reports the early-warning classification performance of the residual stress score for predicting whether a drawdown onset occurs within the next  $H = 21$  trading days. The ROC curve summarizes the trade-off between true- and false-positive rates across score thresholds, while the precision–recall curve is more informative under class imbalance (unconditional positive-label rate  $\bar{y} = \Pr(y_t^{(H)} = 1) \approx 13\%$ ; see Eq. (18)).

Figure 13 reports an event-study view of the residual stress score around drawdown onsets. For each onset, we align days by event time  $\tau$  (with  $\tau = 0$  at the onset) and compute the cross-event mean stress across the  $\pm 30$  trading-day window. The shaded band summarizes dispersion across events.



**Figure 12.** Early-warning performance ( $H = 21$ ): ROC curve (left) and precision–recall curve (right) for the residual stress score.



**Figure 13.** Event-study of residual stress around drawdown onsets ( $\pm 30$  trading days;  $\tau = 0$  at onset).

Table 4 reports early-warning performance for predicting drawdown onsets within the next  $H = 21$  trading days. Panel A summarizes standard classification metrics. In the baseline sample, the SPY-volatility benchmark achieves higher ROC-AUC and PR-AUC than residual stress, indicating stronger average discrimination. Panel B evaluates whether residual stress contains incremental information beyond volatility by conditioning onset probabilities on the joint stress/volatility regimes. Conditional on low volatility, high residual stress increases the onset probability from 7.78% (low-stress/low-vol) to 17.02% (high-stress/low-vol), an increase of approximately  $2.2\times$ . The joint high-stress/high-volatility regime exhibits the highest onset probability (36.0%). Because the high-stress regimes occur infrequently (e.g.,  $N = 47$  and  $N = 25$  days), these conditional estimates should be interpreted with sampling uncertainty. Table 4 indicates that realized volatility outperforms residual stress as a standalone early-warning score in our baseline run. We therefore conduct an *incremental-value test* to assess whether residual stress provides complementary information beyond volatility when both signals are combined under the same leakage-safe timing protocol. This exercise is interpreted as a diagnostic: a meaningful improvement in PR-AUC (or an increase in recall at comparable precision) would support the view that residual stress captures cross-sectional dislocation not already summarized by time-series volatility.

**Table 4.** Early-warning performance and volatility-complementarity diagnostics ( $H = 21$ ,  $\delta = -10\%$ ). Panel A reports aggregate early-warning metrics. Panel B reports regime-conditional onset probabilities under the joint residual-stress and volatility regimes.

**Panel A: Early-warning metrics (higher is better).**

Signal	ROC-AUC	PR-AUC	Precision	Recall	F1	Event rate
Residual stress	0.604	0.162	0.236	0.121	0.160	0.133
SPY volatility	0.685	0.186	0.297	0.386	0.335	0.133

**Panel B: Conditional onset probability by joint regimes.**

Regime (Stress, Vol)	$N$ days	$\mathbb{P}(\text{onset within } H)$
Low, Low (Q00)	1003	0.0778
High, Low (Q10)	47	0.1702
Low, High (Q01)	157	0.2866
High, High(Q11)	25	0.3600

**Panel C: Event-level overlap within a 63-day lookback window.**

Category	$N$ events	Share
Both	8	0.571
Stress only	1	0.071
Vol only	0	0.000
Neither	5	0.357

**Panel D: Lead-time to onset (paired events; lookback = 63).**

Metric	Value
$n_{\text{paired}}$	8
Median lead (stress)	57.5
Median lead (vol)	60.0
Mean lead (stress)	51.75
Mean lead (vol)	57.75
Pr(stress earlier)	0.375

#### 4.7. Complementarity Beyond Volatility: Joint-Regime and Overlap Diagnostics

Aggregate early-warning metrics in Table 4 (Panel A) show that realized volatility outperforms residual stress as a standalone score in the baseline sample. We therefore evaluate whether residual stress provides *complementary* information beyond volatility using diagnostic tests that preserve the same leakage-safe timing protocol.

##### Joint-Regime Conditioning

Table 4 (Panel B) reports the conditional probability of a drawdown onset within the next  $H = 21$  trading days under the joint regimes defined by residual stress and SPY volatility. Conditional on low

volatility, high residual stress increases the onset probability from 7.78% (low-stress/low-vol) to 17.02% (high-stress/low-vol), an increase of approximately  $2.2\times$ . The joint high-stress/high-volatility regime exhibits the highest onset probability (36.0%). Because the high-stress regimes occur infrequently (e.g.,  $N = 47$  and  $N = 25$  days), these conditional estimates should be interpreted with sampling uncertainty.

#### Event-Level Overlap and Lead-Time

To assess whether residual stress captures events not flagged by volatility thresholding, we compute event-level overlap within a fixed lookback window (63 trading days prior to each onset). In the baseline sample, residual stress uniquely flags one onset that is not preceded by a volatility alarm, while volatility does not uniquely flag any onset; both alarms occur before 57% of onsets, and neither alarm occurs before 36% of onsets (Table 4, Panel C). Lead-time diagnostics indicate that, among paired events where both alarms occur, stress and volatility have similar lead time distributions (median lead 57.5 vs. 60 trading days; Table 4, Panel D), suggesting that the primary incremental value of residual stress in this setting is risk stratification conditional on volatility rather than systematically earlier triggering.

#### Exploratory Combined-Score Model

For completeness, we also consider an exploratory two-signal combination using a walk-forward logistic regression re-estimated using information available only through  $t-1$  so that predictions at date  $t$  remain leakage-safe:

$$\Pr(y_t^{(H)} = 1 \mid \text{Stress}_t, \text{Vol}_t) = \sigma(\beta_0 + \beta_1 \text{Stress}_t + \beta_2 \text{Vol}_t), \quad (27)$$

where  $y_t^{(H)}$  indicates whether a drawdown onset occurs within the next  $H$  trading days and  $\sigma(\cdot)$  is the logistic link. In the baseline sample, this combined specification does not improve early-warning performance relative to the volatility benchmark and appears sensitive to modeling choices under walk-forward estimation. We therefore report full details in Appendix A.1 and treat the combined-model exercise as exploratory rather than a primary contribution.

#### 4.8. Sensitivity and Robustness Checks

This study reports a compact sensitivity analysis over key design parameters ( $K$ ,  $q$ ,  $W_{\text{thr}}$ ,  $\delta$ , and  $H$ ) while preserving the same leakage-safe timing protocol throughout. Across the specifications in Table 5, the qualitative conclusions remain unchanged: residual stress contains early-warning information and clusters around drawdown onsets, while the volatility baseline remains highly competitive in predictive performance.

**Table 5.** Sensitivity analysis (leakage-safe). Early-warning metrics for residual stress and the volatility baseline across alternative parameter choices.

Spec	$K$	$q$	$W_{\text{thr}}$	$\delta$	$H$	Residual stress		Volatility	
						ROC-AUC	PR-AUC	ROC-AUC	PR-AUC
Baseline	2	0.95	252	-0.10	21	0.6025	0.1568	0.6917	0.1857
K=1	1	0.95	252	-0.10	21	0.6271	0.1791	0.6849	0.1862
K=3	3	0.95	252	-0.10	21	0.5624	0.1435	0.6849	0.1862
$q = 0.90$	2	0.90	252	-0.10	21	0.6044	0.1622	0.6849	0.1862
$\delta = -0.05$	2	0.95	252	-0.05	21	0.4715	0.2011	0.4583	0.2104
$H = 63$	2	0.95	252	-0.10	63	0.5639	0.3242	0.7017	0.5036

## 5. Discussion

This study proposes a leakage-safe PCA–APT residual stress index constructed from the cross-sectional dispersion of sector excess-return residuals and evaluates whether the resulting stress regimes contain early-warning information for SPY drawdowns. The empirical patterns in Figures 3 and 4 suggest that large residual-stress spikes tend to cluster around drawdown onsets, supporting the interpretation that unusually high idiosyncratic dispersion across sectors can coincide with (and sometimes precede) market stress episodes. This is consistent with a market microstructure intuition: during risk-off transitions, sector-level repricing and rotation can become less synchronized with the low-dimensional common-factor structure, thereby increasing the magnitude of residual dislocations even after removing the market component via excess returns.

### 5.1. Economic Interpretation of the PCA Factor Space

The PCA results indicate a clear low-dimensional structure in sector excess returns. The scree plot and variance decomposition (Table 1, Figure 1) show that the first two components capture a large fraction of common variation, motivating the parsimonious choice  $K = 2$  for factor extraction. The loading heatmap (Figure 2) and top-loading summary (Table 2) further indicate that a subset of sectors contributes disproportionately to the leading components, which aligns with the notion that sector cross-sections are driven by a small number of latent macro/industry forces. Importantly, the stress index is not defined from the factors themselves, but from the residual component orthogonal to the factor space; hence, the proposed signal is intended to capture “abnormal” cross-sectional dislocation beyond what is explained by dominant common variation.

### 5.2. Residual Stress as a Regime Indicator Rather than a Linear Return Predictor

Figures 9 and 10 suggest that the contemporaneous residual stress score has only a modest average linear association with forward SPY returns, and that any stress–return linkage is time-varying. This has two implications. First, the stress index should be interpreted primarily as a *regime* indicator (i.e., a measure of tail-risk conditions) rather than a stable return-forecasting factor. Second, evaluation should emphasize event-based and classification-style metrics (e.g., ROC-AUC, precision/recall for drawdown-onset warnings) in addition to mean-return regressions, because the economic value may lie in identifying periods when risk management is particularly valuable rather than predicting the sign of average returns in normal times.

### 5.3. Implications for Risk Management Overlays

The stress-managed SPY overlays provide a simple demonstration of how a residual-dispersion signal can be operationalized for risk control. In the reported run (Table 3; Figures 5 and 6), both overlay variants track the baseline closely, which indicates that the chosen exposure reduction and threshold level yield a conservative effect over this sample. This outcome can be desirable if the goal is a low-tracking-error overlay, but it also highlights a practical trade-off: stronger protection typically requires either (i) a more aggressive de-risking rule (larger exposure reduction), (ii) a lower stress threshold (more frequent activation), or (iii) a continuous scaling rule that maps stress magnitude to exposure rather than using a binary label. A natural extension is to introduce hysteresis or smoothing (e.g., require stress to remain elevated for  $m$  days before switching, or use a rolling mean of stress) to reduce turnover and improve net performance under transaction costs. Additionally, instead of a fixed haircut, one can scale exposure by normalized stress (e.g., z-scored relative to the rolling window) to better differentiate moderate from extreme stress episodes.

### 5.4. Why the Residual-Based Long–Short Can Underperform

The residual-ranked APT long–short strategy underperforms in the example summary (negative annualized return and deeper drawdown in Table 3), while Figure 11 shows nontrivial turnover. This is not surprising in a daily-rebalanced cross-sectional setting: residual signals can be noisy, and ranking-based portfolios often require either (i) stronger signal-to-noise (e.g., additional conditioning

variables), (ii) longer holding periods, or (iii) explicit constraints and regularization to control turnover. Even if residual dispersion is informative about *market stress*, it does not necessarily imply that individual residuals provide reliable short-horizon cross-sectional alpha after costs. Practically, several refinements could improve realizability: using multi-day averaged residual signals, volatility scaling, sector-weight caps, turnover penalties, and/or less frequent rebalancing. These results therefore support viewing the L/S strategy as an auxiliary diagnostic of residual informativeness, while the central contribution remains the stress-regime identification and its risk-management application.

### 5.5. Robustness Considerations and Limitations

Section 4 reports a compact sensitivity analysis over key design parameters ( $K$ ,  $q$ ,  $W_{\text{thr}}$ ,  $\delta$ , and  $H$ ) under the same leakage-safe timing protocol (Table 5). Nonetheless, several limitations remain and additional robustness checks are warranted in broader samples. First, the factor dimension  $K$  governs what is treated as “common” versus “idiosyncratic”; increasing  $K$  mechanically reduces residual magnitude and may change the timing of stress signals. Second, the rolling-window length  $W$  and quantile level  $q$  control the out-of-sample threshold; shorter windows adapt faster but may be noisy, while higher quantiles produce rarer, more extreme stress labels. Third, the drawdown threshold  $\delta$  and the early-warning horizon (e.g., 21 trading days) define the event being predicted; varying these parameters helps determine whether the signal is specific to large, slow-moving drawdowns or also captures milder corrections. In addition, the analysis focuses on a sector-ETF cross-section and a single market proxy (SPY); generalization to other equity universes, international indices, and alternative asset classes is an important direction for future work. Finally, while the implementation is designed to be leakage-safe via rolling, train-only thresholding and lagged signals, results can still be sample-dependent, and the use of public data sources introduces potential issues related to survivorship, symbol changes, and data revisions. Accordingly, future work should include broader samples, additional validation windows, and alternative data vendors where feasible.

### 5.6. Future Research Directions

Two extensions are especially promising. First, combining residual stress with complementary stress proxies (e.g., implied volatility, credit spreads, or liquidity measures) may improve drawdown early-warning performance and reduce reliance on a single signal. Second, replacing the binary threshold rule with a probabilistic model (e.g., logistic regression or survival-style hazard modeling for drawdown onset) could map the stress index into an interpretable drawdown-risk probability and support decision-theoretic overlay design. More broadly, connecting cross-sectional residual dispersion to macro-financial mechanisms (risk premia, funding liquidity, or sector rotation) would strengthen the economic interpretation and improve external validity.

## 6. Conclusions

This paper develops a leakage-safe PCA–APT residual stress index for equity-market drawdown monitoring using a cross-section of U.S. sector ETFs. The key idea is to extract a low-dimensional factor structure from sector *excess* returns (relative to SPY) via PCA and to define market “stress” as the cross-sectional dispersion of the factor *residuals*. A rolling, out-of-sample quantile rule is then used to convert the continuous stress score into a regime label, enabling drawdown-onset analysis and risk-management backtests without look-ahead bias.

Empirically, the residual stress index exhibits pronounced clustering around major SPY drawdown episodes, and the regime shading aligns visually with periods of deteriorating market conditions. While the stress score has only modest average linear association with forward SPY returns, its primary value lies in identifying high-risk *states* rather than serving as a stable return-forecasting factor. In backtests, a stress-managed SPY overlay illustrates how the signal can be translated into a practical de-risking rule with controlled tracking error, whereas a residual-ranked sector long–short strategy highlights the challenges of converting noisy daily residual information into net cross-sectional alpha after transaction costs.

Overall, the results support the residual-dispersion channel as a useful and interpretable lens for market stress monitoring and drawdown risk control. Future work should broaden the asset universe, stress-test sensitivity to factor dimension and threshold design, and evaluate event-based predictive metrics more formally. Promising extensions include probabilistic drawdown-hazard models that map the stress index into a calibrated risk probability, multi-signal ensembles that combine residual stress with volatility and credit/liquidity indicators, and overlay designs that scale exposure continuously with stress magnitude while explicitly penalizing turnover. These directions can improve robustness and strengthen the case for residual-stress indicators as tools for systematic risk management in equity markets.

**Author Contributions:** Conceptualization, T.L.; methodology, T.L.; software, T.L.; validation, T.L.; formal analysis, T.L.; investigation, T.L.; data curation, T.L.; writing—original draft, T.L.; writing—review & editing, T.L.; visualization, T.L. The author has read and agreed to the published version of the manuscript.

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

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** Data were obtained from Yahoo Finance. The datasets analyzed during the current study are publicly available from these sources.

**Acknowledgments:** The author used AI-assisted tools for language editing and formatting. The author reviewed and takes responsibility for the content.

**Conflicts of Interest:** The author declares no conflicts of interest.

## Appendix A

### Appendix A.1. Combined-Model Incremental-Value Test

This appendix reports the baseline combined-model results for the walk-forward logistic specification in Eq. (27). The combined model underperforms the standalone volatility benchmark in this sample and is included for transparency.

**Table A1.** Incremental-value test: early-warning performance of the combined (stress+volatility) model.

Model	H	N	Event rate	ROC-AUC	PR-AUC	Precision	Recall	F1
Combined (Stress + Vol)	21	1543	0.130092	0.408781	0.051975	0.081633	0.070175	0.075472

## References

- Pandey, P.N. ETF Risk–Return and Price Discovery in Crisis Regimes: Evidence from Indian Markets. *SSRN* **2025**, 5899923.
- Ross, S.A. The arbitrage theory of capital asset pricing. *Journal of Economic Theory* **1976**, *13*(3), 341–360.
- Jolliffe, I.T. *Principal Component Analysis*, 2nd ed.; Springer: New York, NY, USA, 2002.
- Campbell, J.Y.; Lettau, M.; Malkiel, B.G.; Xu, Y. Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *The Journal of Finance* **2001**, *56*(1), 1–43.
- Christie, W.G.; Huang, R.D. Following the Pied Piper: Do individual returns herd around the market? *Financial Analysts Journal* **1995**, *51*(4), 31–37.
- Chang, E.C.; Cheng, J.W.; Khorana, A. An examination of herd behavior in equity markets: An international perspective. *Journal of Banking & Finance* **2000**, *24*(10), 1651–1679.
- Moreira, A.; Muir, T. Volatility-managed portfolios. *The Journal of Finance* **2017**, *72*(4), 1611–1644.
- Barroso, P.; Santa-Clara, P. Momentum has its moments. *Journal of Financial Economics* **2015**, *116*(1), 111–120.
- Hanley, J.A.; McNeil, B.J. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* **1982**, *143*(1), 29–36.
- Davis, J.; Goadrich, M. The relationship between Precision–Recall and ROC curves. In *Proceedings of the 23rd International Conference on Machine Learning (ICML)*; 2006; pp. 233–240.

11. Ciciretti, V.; Nandy, M.; Pallotta, A.; Lodh, S.; Senyo, P.K.; Kartasova, J. An early-warning risk signals framework to capture systematic risk in financial markets. *Quantitative Finance* **2025**, *25*, 757–771.
12. Febrian, E.; Herwany, A. CAPM and APT validation test before, during, and after financial crisis in emerging market: Evidence from Indonesia. *SSRN* **2010**.
13. Koutoulas, G.; Kryzanowski, L. Integration or segmentation of the Canadian stock market: Evidence based on the APT. *Canadian Journal of Economics* **1994**, 329–351.
14. Baghdadabad, M.R.T.; Glabadianidis, P. An extensile method on the arbitrage pricing theory based on downside risk (D-APT). *International Journal of Managerial Finance* **2014**, *10*, 54–72.
15. Yang, L.; Rea, W.; Rea, A. Financial insights from the last few components of a stock market PCA. *International Journal of Financial Studies* **2017**, *5*, 15.
16. Neves-Silva, R.; et al. Complex principal component analysis of dynamic correlations in financial markets. In *Intelligent Decision Technologies: Proceedings of the 5th KES International Conference on Intelligent Decision Technologies (KES-IDT 2013)*; Volume 255; p. 111; 2013.
17. Quiroga-Juárez, C.A.; Villalobos-Escobedo, A. Analysis of stock market behavior of the major financial exchanges worldwide using multivariate analysis (principal component analysis PCA) for the period 2011 to 2014. Working Paper, Instituto Tecnológico Metropolitano (ITM), **2025**.
18. Geboers, H.; Depaire, B.; Annaert, J. A review on drawdown risk measures and their implications for risk management. *Journal of Economic Surveys* **2023**, *37*, 865–889.
19. Acharya, V.; Steffen, S. Stress tests for banks as liquidity insurers in a time of COVID. *VoxEU.org* **2020** (March).