Dynamic Fear Hedging: Using the “Fear Gauge” to Mitigate Risk

Ahmer Nadeem Khan

Erdos Institute: Quant Finance Bootcamp

November 7, 2025

Table of Contents

Table of Contents

Yahoo Finance Scraping

- `yfinance` tends to unreliable for large or repeated queries.
- Wrote **custom scraping scripts** using direct Yahoo Finance URLs to collect historical data for any ticker, frequency (daily, weekly, monthly), and field (open, close, high, low, volume etc.).
- Browser automation and page rendering using `Playwright`.
- Parsed website HTML tables using `Beautiful Soup`.
- Implemented `asynchronous (async)` execution for efficient automation.
- Built in failure-handling and retry logic.
- Refer to `01_data_collection.ipynb`.

Data Collected

We collect the data for the following instruments (tickers):

1. S&P 500 (GSPC); used as a proxy for S&P 500 INDEX (SPX)
2. CBOE Volatility Index (VIX)
3. iPath Series B S&P 500 VIX Short-Term Futures ETN (VXX)

We get daily Year-to-date (YTD) data i.e. January 02 2025 - November 06 2025 (total trading days = 213) including open, close, high, low, adj. close (where applicable/available) and volume (where applicable/available). The data is collected and saved in a CSV format and can be read/written easily with pandas functionality.

Note that markets are closed on January 01.

Additional Functionality (/scripts)

The /scripts folder contains several auxiliary routines used in earlier large-scale data collection projects. Each script outputs CSV files in the current directory.

- **scrape_oneStock:** Fetches **weekly adjusted close prices** for a single ticker (default: AAPL) over the past two years.
- **scrape_allStocks:** Collects weekly adjusted close data for an entire ticker list (e.g. S&P 500). Requires `sp500_tickers.csv`; retries up to 3 times per ticker and logs failures in `failed_tickers.csv`.
- **scrape_S&P_ordered:** Scrapes and outputs the list of S&P 500 constituents **ordered by market cap**. Uses data from *Slickcharts* (manual browser mimic, not Playwright).
- **combine_CSVs:** Merges two CSV files into one — used to combine results from multiple scraping rounds.

Each script can be customized for alternative frequencies, fields (beyond adj. close), or data sources. See notebook for detailed explanations on how to use these scripts.

Table of Contents

Visualization

- We present simple charts for our initial data, and then an overlay chart.
- Refer to `02_visualization.ipynb`.

S&P 500 Index – YTD Performance

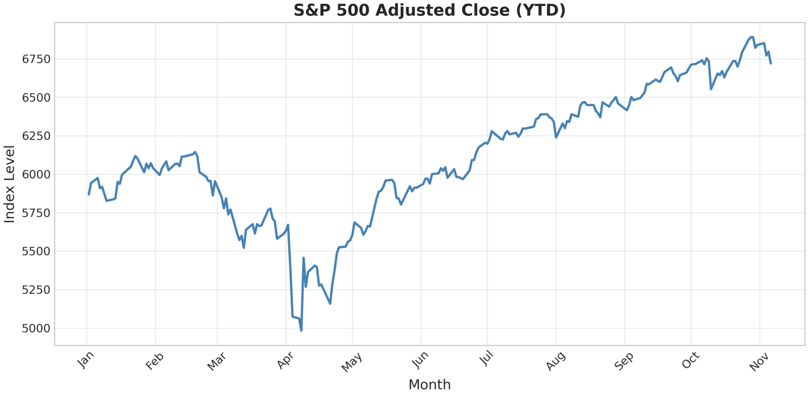

Figure: S&P 500 daily closing prices in 2025 YTD.

Introduction to the VIX

- The **CBOE Volatility Index (VIX)** measures the market's expectation of near-term (30 day) volatility implied by S&P 500 option prices.
- The calculation takes as input the market prices of SPX options and SPXW options as well as U.S. Treasury yield curve rates.
- Often called the **"fear index"**, it tends to rise sharply during periods of market stress or uncertainty.
- VIX reflects **expected (implied) volatility**, not realized volatility — it is forward-looking.
- The index is mean-reverting and exhibits strong **negative correlation** with the S&P 500.

Calculating the VIX

The VIX represents the square root of the risk-neutral expected variance of the S&P 500 return over the next 30 days.

$$\sigma^2 = \frac{2}{T} \sum_i \frac{\Delta K_i}{K_i^2} e^{RT} Q(K_i) - \frac{1}{T} \left[\frac{F}{K_0} - 1 \right]^2$$

$$\text{VIX Index} = \sigma \times 100$$

where:

- T — Time to expiration (in years)
- F — Option-implied forward price
- K_0 — First strike below or equal to the forward index level F
- K_i — Strike of the i^{th} out-of-the-money option
- $\Delta K_i = \dfrac{K_{i+1} - K_{i-1}}{2}$ — Interval between strike prices
- $Q(K_i)$ — Midpoint of bid-ask spread for each option with strike K_i
- R — Risk-free interest rate to expiration

VIX Index – YTD Performance

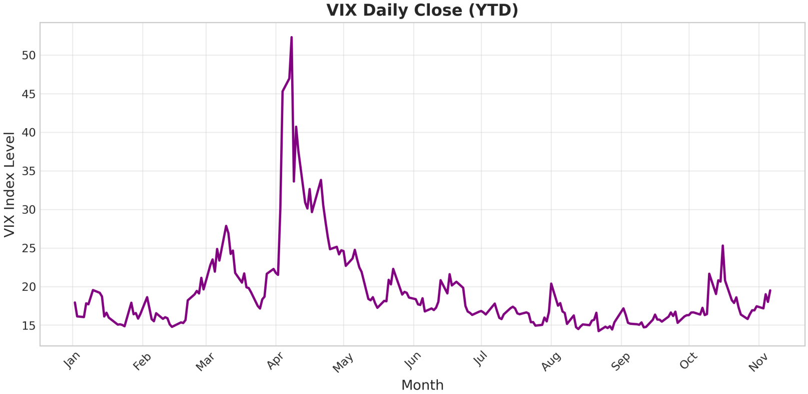

Figure: Daily closing levels of the CBOE Volatility Index (VIX) during 2025 YTD.

Distribution of Daily Returns (2025 YTD)

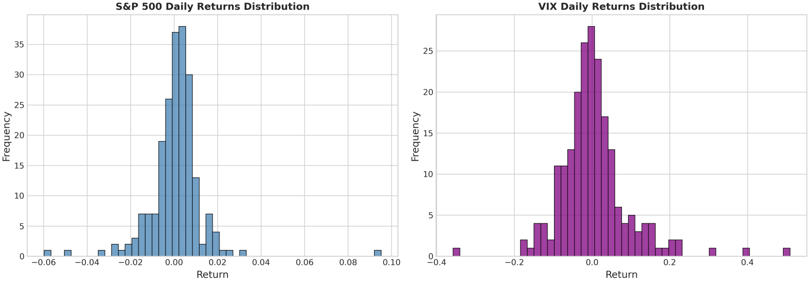

Figure: Comparison of daily return distributions for VIX and S&P 500. Note the axis resolutions.

SPX vs VIX – Dual Axis View (YTD)

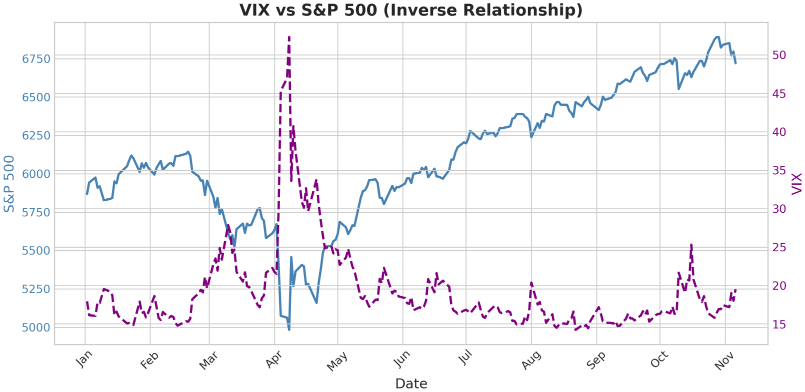

Figure: Inverse co-movement between S&P 500 and VIX levels throughout the year. Major VIX spikes align with SPX drawdowns.

Event Study

We identify **market fear spikes** in the VIX using both level- and jump-based criteria, combining statistical thresholds with temporal filtering.

1. **Level-based rule:**
   - A day is flagged if VIX $\geq 95^{\text{th}}$ percentile of all YTD VIX closes.
2. **Jump-based rule:**
   - A day is flagged if the daily VIX change ΔVIX $\geq \mu + 2\sigma$, where μ and σ are the mean and standard deviation of daily changes.
3. **Cooldown filter:**
   - Spikes within ± 3 trading days are merged — only the largest VIX close in each cluster is retained.
4. **Enrichment:**
   - Each detected spike is tagged with the same-day S&P 500 return and ΔVIX.

Detected VIX Spikes

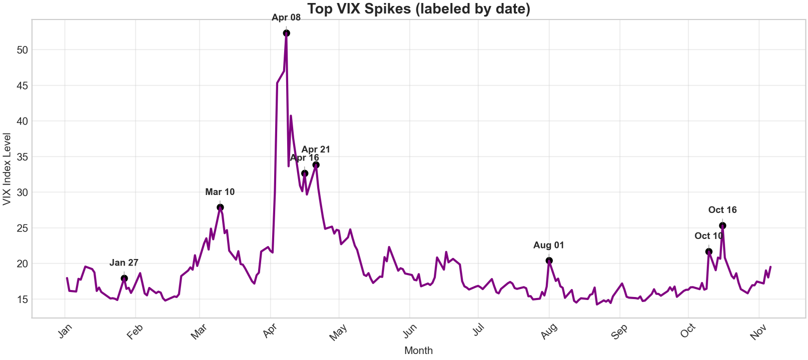

Figure: Annotated spikes correspond to top jump/level spike detections

Market Events

- Event selection combines quantitative filtering (large VIX daily % jumps) with qualitative context (news verification via Bloomberg, CNN, etc.).
- This was done manually using expert sources.
- The annotations were added using matplotlib functionality, and sources were also added to the image.
- They are also available in the annotations CSV on GitHub.
- Refer to `03_event_study.ipynb`.

VIX Annotated with Key Market Events (2025 YTD)

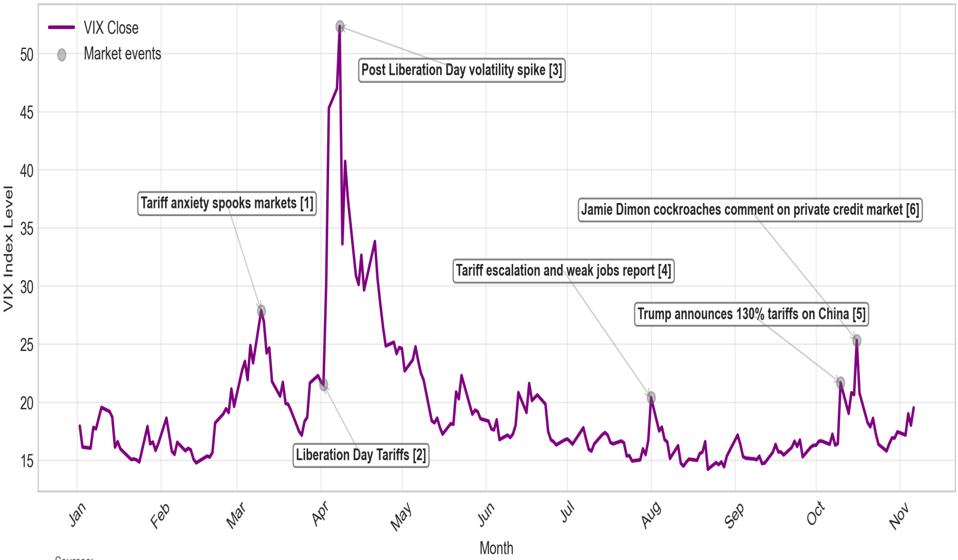

Sources:

[1] https://www.bloomberg.com/news/articles/2025-03-10/vix-nears-30-first-time-since-august-shock-on-tariff-anxiety

[2] https://en.wikipedia.org/wiki/Liberation_Day_tariffs

[3] https://www.kitco.com/news/off-the-wire/2025-04-07/stocks-plunge-tariff-turmoil-vix-fear-gauge-spikes

[4] https://sweetvolatility.com/volatility-market-recap-july-28-august-2-2025

[5] https://www.cnn.com/2025/10/10/economy/trump-china-tariff-threats-economy

[6] https://www.cnn.com/2025/10/16/business/jamie-dimon-us-economy-cockroaches

Table of Contents

Initial Analysis

- We investigate the relationship between changes in the S&P 500 index and changes in the VIX.
- We fit a regression line using `statsmodels`, investigate the correlation (R^2 value) and visualize the relationship.
- We also report on the residuals and fitting for completeness.
- This will motivate our portfolio/hedging goals.
- Refer to `04_analysis.ipynb`.

Regression: S&P 500 vs VIX Daily Returns

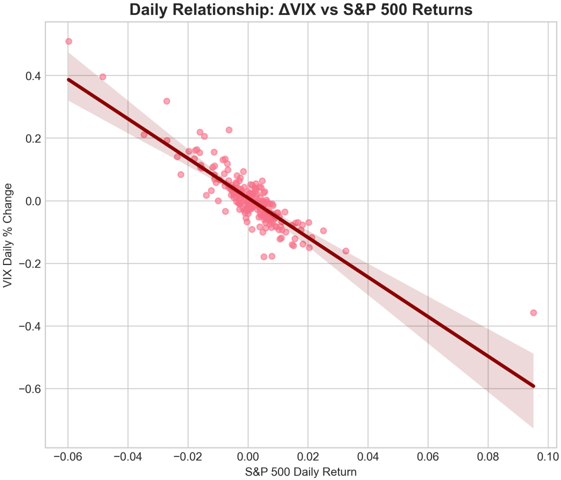

Figure: Regression of daily VIX returns on S&P 500 returns shows a strong negative slope, confirming the inverse relationship between equity and volatility.

Regression Summary: SPX vs VIX Daily Returns

We estimate the following regression model to capture the relationship between equity and volatility returns:

$$y = -6.3183x + 0.0088 + \varepsilon$$

$$R^2 = 0.753$$

where:

- y = VIX daily return ($dVIX$)
- x = S&P 500 daily return ($dSPX$)
- ε = residual term

Results

Interpretation:

- The slope of -6.32 indicates a strong negative relationship between equity and volatility.
- $R^2 = 0.75$ implies that about 75% of daily VIX variation is explained by S&P 500 returns.
- Historically, longer-term values of R^2 are around 0.3–0.4; our YTD result shows unusually high explanatory power, and is encouraging for further analysis.

OLS Regression based on 212 daily observations (2025 YTD).

Regression Diagnostics

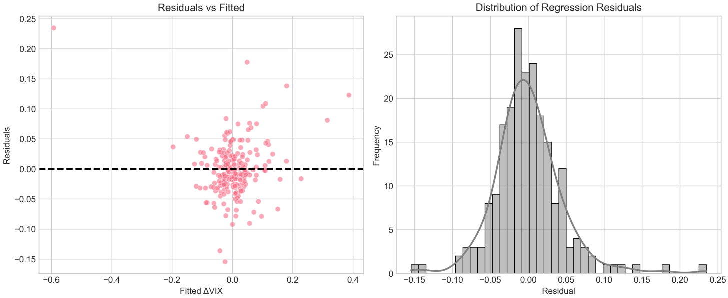

Figure: Residual analysis of the SPX–VIX regression indicates homoskedastic residuals and a stable negative relationship.

Table of Contents

Hedging with VIX

- We investigate how to hedge exposure to equities using the VIX, i.e. hedge a long position in large cap stocks.
- We use an ETF that has very high correlation to the volatility index as a proxy investment and provide evidence for its hedging ability.
- We use the S&P 500 index as a proxy for SPX since they have near-identical returns due to the high liquidity and low tracking error.
- Refer to `05_hedging_portfolio.ipynb`

VIX Exposure Instrument

Exposure to the **VIX Index** is generally difficult for investors to obtain directly. It can be achieved through:

- **VIX options**,
- **VIX futures** (provided by CBOE), or
- **ETFs / ETNs** that track VIX options or futures.

In this study, we use the short-term ETF that tracks the **S&P 500 VIX Short-Term Futures Index** — namely the **iPath Series B S&P 500 VIX Short-Term Futures ETN (VXX)**.

This product is **not appropriate for long-term buy-and-hold hedged portfolios**, as it tracks futures with an average one-month maturity and resets its exposure daily.

VXX was launched on January 19, 2018 and is issued by iPath. It was chosen primarily due to the ease of data access via Yahoo Finance for our scraping routines.

Correlation Matrix

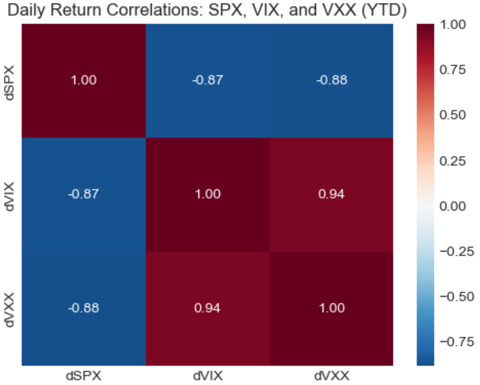

Figure: Correlation heatmap of daily returns between the S&P 500, VIX, and VXX.

$$\mathrm{Corr}(SPX, VIX) = -0.87$$
$$\mathrm{Corr}(SPX, VXX) = -0.88$$
$$\mathrm{Corr}(VIX, VXX) = 0.94$$

Rolling Correlations

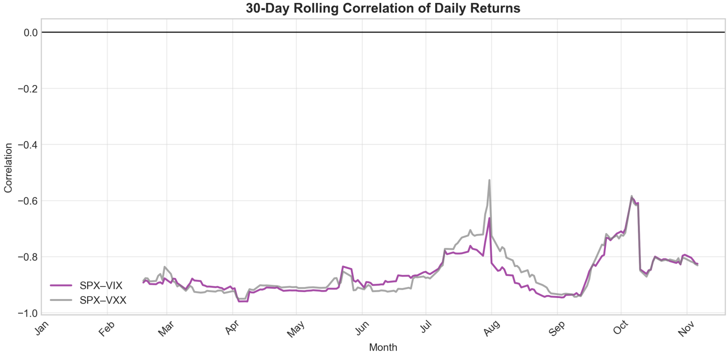

Figure: Rolling 30-day correlation of S&P 500 returns with VIX and VXX, showing tracking of the index.

SPX vs VIX/VXX

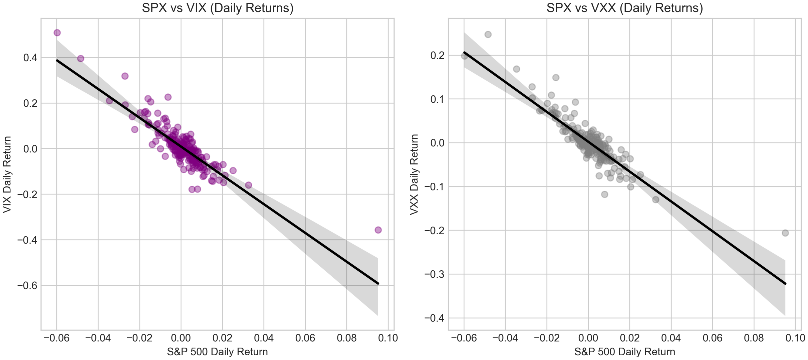

Figure: Scatter plots with regression lines showing the negative relationship between S&P 500 and volatility indices.

Static Hedging

We compute the **Minimum Variance Portfolio** P comprised of:

- **SPX (B):** Benchmark placeholder for the market holdings (M)
- **VXX:** VIX futures ETF (volatility exposure)

The portfolio returns are given by:

$$R_P = R_{SPX} + \beta \times R_{VXX}$$

and the optimal weight β^* (optimal hedge ratio) is given by:

$$\beta^* = -\frac{\mathrm{Cov}(R_{SPX}, R_{VXX})}{\mathrm{Var}(R_{VXX})}$$

Here, β^* minimizes portfolio variance and represents the static hedge weight for VXX.

Cumulative Returns: SPX vs Static Hedge

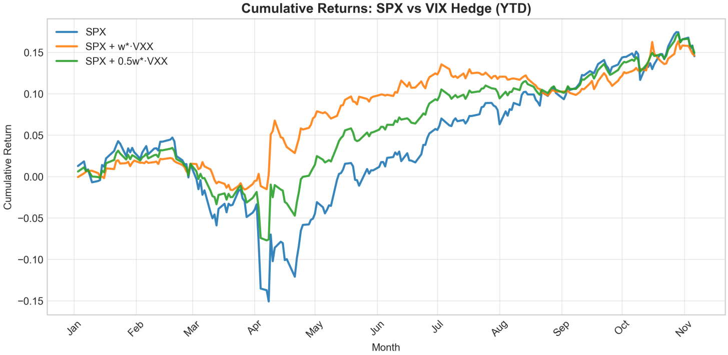

Figure: Comparison of cumulative returns for the S&P 500 and static VIX-hedged portfolios.

Basic Performance Metrics

- **Annualized Volatility:** Measures annualized portfolio risk.

$$\sigma_{\text{annualized}} = \sigma_{\text{daily}} \times \sqrt{\text{trading days}}$$

- **Sharpe Ratio:** Measures excess return per unit of risk.

$$SR_p = \frac{\mathbb{E}[R_p - R_f]}{\sigma_p} \quad (\text{set } R_f = 0 \text{ for simplicity})$$

- **Maximum Drawdown:** Largest observed loss from a peak to a trough.

$$\text{MaxDD} = \min_t \left(\frac{V_t - \max_{s \le t} V_s}{\max_{s \le t} V_s} \right)$$

A MaxDD of -0.25 (–25%) means the portfolio fell 25% from its peak before recovery.

Performance Report

Table: Final Cumulative Returns % (YTD as of Nov 6, 2025)

Date	**SPX**	**Optimal Hedge**	**Milder Hedge** ($\beta = \frac{1}{2}\beta^*$)
2025-11-06	14.5	14.6	14.9

Table: Portfolio Performance Metrics (YTD)

Portfolio	**Ann. Vol (%)**	**Sharpe**	**Max DD (%)**
SPX	19.76	0.91	-18.90
Optimal Hedge (β^*)	9.33	1.78	-3.78
Milder Hedge ($\beta = \frac{1}{2}\beta^*$)	12.76	1.36	-10.79

Comparable overall return (consistent with risk-premium)
Risk more than halved, and Sharpe ratio almost doubled
Significant reduction in downside risk

Results

These results are very promising and consistent with the predictions of **Portfolio Theory**!

- The hedged portfolio tracks the underlying asset (**SPX**) very well.
- In calm periods, the hedge slightly underperforms SPX (due to the risk premium).
- In downturns, the hedge performs very well, even providing short-term upside value — this arises from the **asymmetry in volatility**: bear markets are more volatile than bull markets.
- The hedge strongly resists large drawdowns in SPX.
- In bullish markets, the hedge also outperforms the index — and by the largest margin — since both the higher volatility and the rising market benefit the position.
- The milder hedge blunts all these effects, offering smoother but smaller protection.

Dynamic Hedging

- We want to hedge dynamically assuming a time-series data, reflecting a real-life trading environment.
- We use the rolling correlations and variances to compute the rolling optimal weight for hedging.
- We lag the weight by 1 day to avoid **look-ahead bias**.
- We use a rolling window of 60 days to compute the weights
- This means our strategy begins after the first 60 days for the pipeline to begin

Dynamic Hedge Formulation

We dynamically adjust the hedge weight w_t to minimize portfolio variance:

$$\min_{w_t} \mathrm{Var}(R_{p,t}) \quad \text{where} \quad R_{p,t} = R_{S,t} + w_t R_{V,t}$$

The optimal hedge weight at each time t is given by:

$$w_t^* = -\frac{\mathrm{Cov}(R_{S,t}, R_{V,t})}{\mathrm{Var}(R_{V,t})}$$

In practice, w_t^* is estimated using a rolling window (e.g., 60 days):

$$\hat{w}_t = -\frac{\widehat{\mathrm{Cov}}_{t-L:t}(R_S, R_V)}{\widehat{\mathrm{Var}}_{t-L:t}(R_V)} \quad \text{for window length } L.$$

Here, $R_{S,t}$ denotes S&P 500 returns and $R_{V,t}$ denotes VIX (or VXX) returns.

Performance Report: Dynamic vs Static

Table: Dynamic Hedge Performance Metrics (YTD)

Portfolio	Annual Volatility (%)	Sharpe Ratio	Max Drawdown (%)
SPX	19.76	0.91	-18.9
Static Hedge	9.33	1.78	-3.78
Dynamic Hedge (60d)	10.40	2.39	-3.80

Note the comparable risk metrics (volatility and max drawdown)

Much higher Sharpe Ratio for the dynamic hedging strategy

$$1.78 \to 2.39$$

Static vs Dynamic Hedging (Event Weeks)

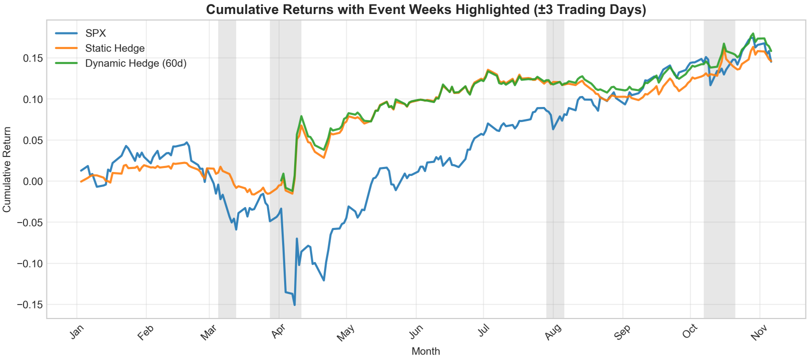

Figure: Static and Dynamic VIX-hedged portfolio performance with event weeks (±3 trading days) shaded.

Event-Week Performance (±3 Trading Days)

Table: Cumulative return over a ±3 trading–day window centered on each event date

Event Date	**SPX (%)**	**Optimal Hedge (%)**	**Dynamic Hedge (%)**
2025-03-10	-4.441	-1.696	
2025-04-02	-11.084	0.025	-1.213
2025-04-08	-5.424	7.274	7.828
2025-08-01	-0.700	-0.107	-0.232
2025-10-10	-1.027	1.612	1.069
2025-10-16	2.790	0.677	1.373

Notes: Blank cell indicates unavailable dynamic series for that window.

Rolling Hedge Weights (w_t) — Dynamic Hedging

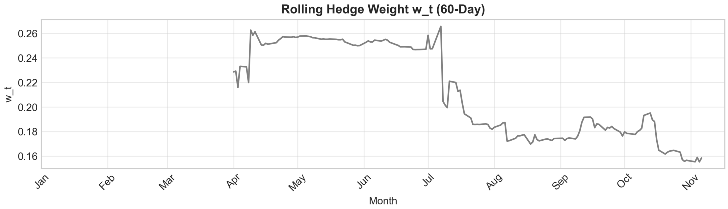

Figure: Estimated rolling hedge ratio (w_t) using a 60-day rolling window. The weight adjusts dynamically in response to changes in correlation and volatility.

Practical Considerations for Dynamic Hedging

- Note that in general, we can observe that the dynamic hedging strategy does better than the static constant weight especially during calmer periods - there is value in re-balancing.
- The results indicate that **dynamic hedging is worth it** — it provides meaningful downside protection and improves overall risk-adjusted returns.
- Regardless, in practice we can **only dynamically hedge**:
  - since the entire dataset is not available *a priori* to compute the optimal static weight.
- In fact, in practice we can only **rebalance at discrete time steps**.
  - This approximates the running dynamic hedge weight function (w_t) shown earlier.
- In reality, hedging frequency is also limited by:
  - convenience, liquidity constraints, and frictions such as transaction costs or roll yield in VIX-based products.

Table of Contents

Conclusions

Given our parameters:

- Real market events/news explains spikes in the VIX
- VIX and its derivatives provide good exposure to volatility
- Exposure to the VIX can be used effectively to hedge a long-position in stocks
- Dynamic hedging is effective and improves hedging characteristics
- Thus, a real-time portfolio may be constructed

Future Considerations

- Sentiment analysis on market news/social media to predict VIX spikes
- Longer time horizons for hedging analysis
- Different rolling windows/parameters etc.
- Simulating VIX using stochastic models
- VIX vs Realized Volatility
- Relationship between VIX and Implied Volatility
- Pricing Options/Futures on the VIX
- Other variance instruments

Table of Contents

References

[1] Yahoo Finance. https://finance.yahoo.com/
[2] CBOE VIX. (2025, September 11). Volatility Index Methodology: CBOE Volatility Index
[2] GitHub Repoistory. (2025, November 07). https://github.com/ahmer-cpu/vix-project