VecViz Analytics Performance As MVO Portfolio Optimization Inputs

Rodger Coyne, Founder of VecViz LLC

2025-11-11

Table of contents

Objectives

1) To determine which VecViz analytic fields, if any, add value as inputs to mean variance portfolio optimization (MVO) relative to trailing 252d return ("Trailing") based inputs.
2) To determine whether portfolios generated by VecViz metric based MVO can outperform a 1/n portfolio and SPY (the SPDR S&P 500 ETF Trust, an ETF that closely tracks the S&P 500 Index).

Methodology Overview

Classical MVO relies on three core inputs: expected returns, risk (volatility/covariance), and diversification (correlation structure). An objective (e.g., maximize expected return) is required, and constraints (e.g., max individual exposure weight, max expected risk) are typically included.

We run MVO for all combinations of VecViz and Trailing MVO inputs using the same constraints and rebalance frequencies, over the same set of tickers, for the same time period.

The strategy universe studied encompasses three distinct model input categories:

1. Trailing based inputs only (S1).
2. Hybrid models that source at least one of the three MVO inputs is Trailing based and at least one is VecViz based (S4,S5, S6,S7, S8, S9, S11, S12, S13).
3. VecViz based inputs only (S10, S14).

The sources of the key return, volatility and correlation inputs to these strategies are detailed further in the table entitled "MVO Input Sources by Strategy", which follows below.

All strategies seek to maximize expected return subject to a maximum expected volatility constraint. Our analysis systematically evaluates these strategies through a rigorous grid search framework which spans multiple constraint configurations including:

1. rebalancing frequencies (10, 21, and 63 days);
2. maximum position weights (3%, 6%, and 10%);
3. max expected volatility constraints (10%, 15%, and 20%).

As implied by the above, each strategy in each objective grouping is ultimately evaluated under 27 different optimization parameter sets (3 rebalance frequencies x 3 maximum position weights x 3 maximum expected volatility constraints). The variation in inputs by strategy and each strategy's objectives are detailed in the performance tables presented in later sections. Each performance table details average performance for each of the 9 constraint and rebalance levels, across all 9 combinations for the other 2 criteria, and then overall, across all 27 parameter sets.

Performance metrics concerning return and risk, and the ratio between them, are standardized and aggregated into a "SummaryZ" score. The average SummaryZ score for all MVO's that each input is part of are then compared to identify the relative value added of each.

MVO Input Sources by Strategy

Here we display the major inputs for each strategy studied. These sources are also conveyed in each strategy's name, with the initials following the underscore denoting input sources for return, volatility, and correlation, in that order. Specifically,

1) T = Trailing 252d;
2) V = VecViz;
3) VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint).

Thus, for example, S8_TVVe is a strategy that uses Trailing for the expected return input, VecViz for the ticker vol, and VecViz's VecEvent similarity measure for correlation between tickers.

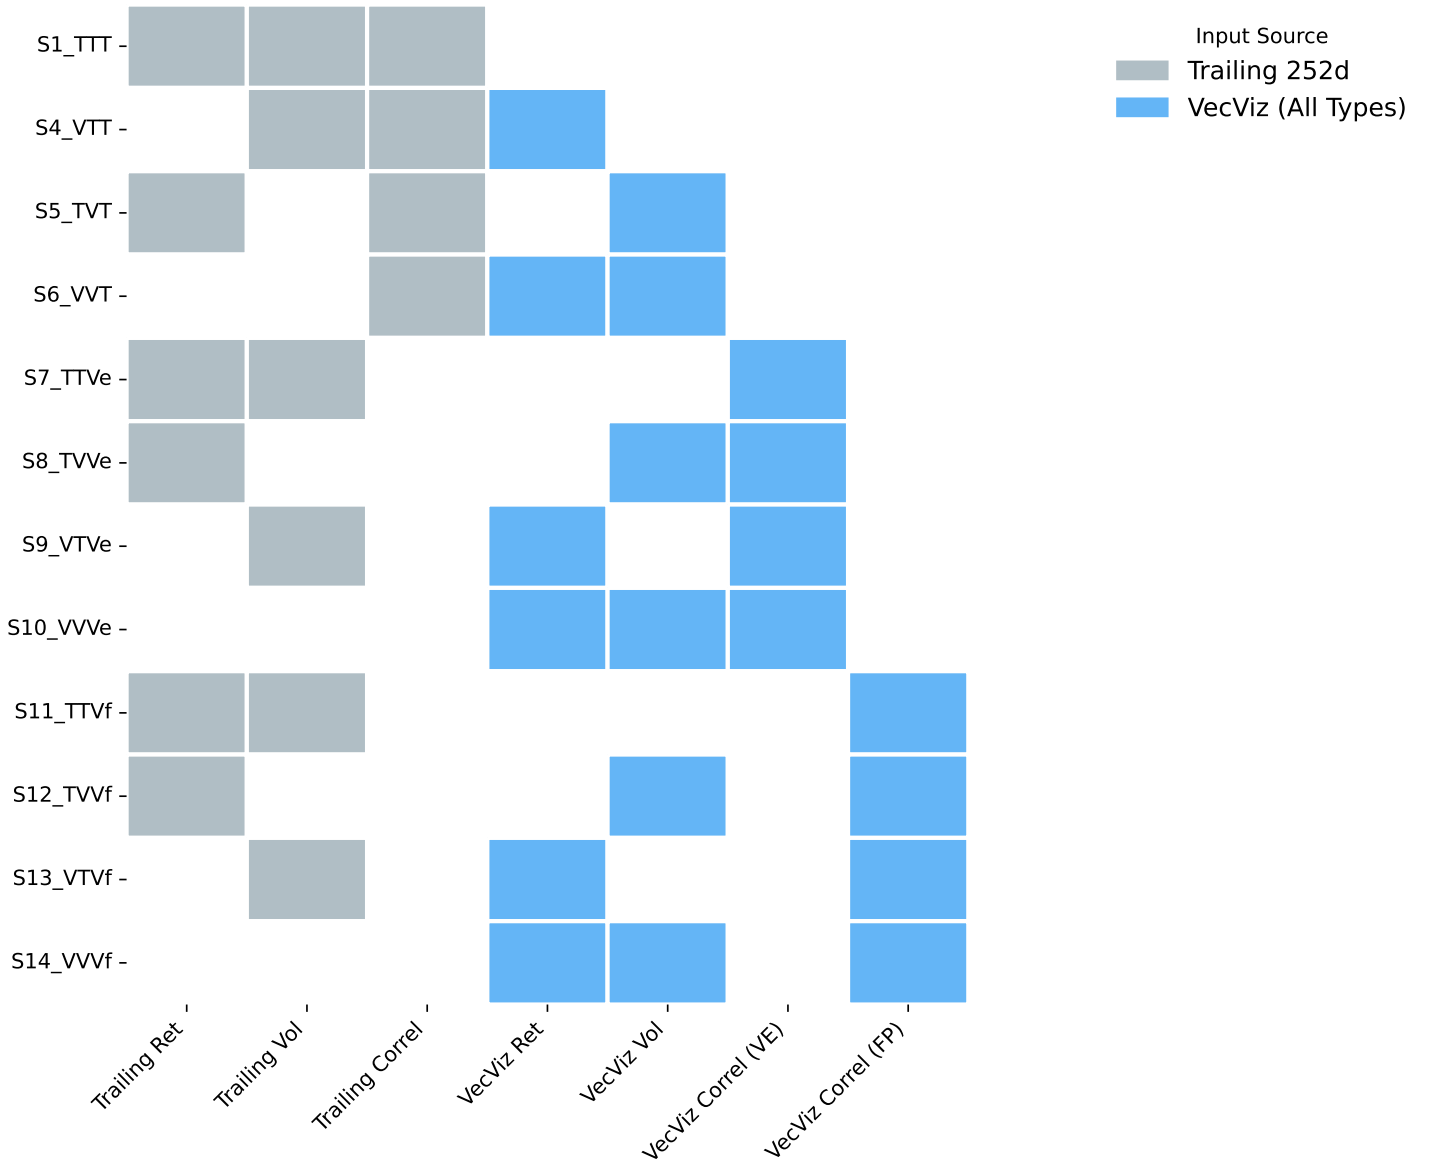

The table reflects balance and consistency with regard to use of VecViz vs. Trailing for each input. For example, there are six that use Trailing for expected return (S1, S5, S7, S8, S11, S12), and six that use VecViz (S4, S6, S9, S10, S13, S14). Likewise there are six that use Trailing 252d for ticker level vol and six that use VecViz for ticker level vol. Finally, there are four that use Trailing 252d for ticker correlation, four that use VecViz's VecEvent

approach, and four that use VecViz's Fingerprint approach.

Note that there is no "S2" or "S3" strategy discussed in this study. Apologies for any associated confusion.

Data

All metrics presented are calculated using daily closing price data sourced from QuoteMedia.

VecViz's MVO inputs are sourced from the inputs and outputs to its Vector Model and V-Score, machine learning models of price probability and forward relative price performance, respectively.

The Vector Model was trained upon ~ 60,000 ticker model dates (TMDs) representing ~550 tickers (including equities, currencies, and commodities) and ~ 120 model dates spanning from 3/9/2002 to 2/3/2021. The Vector Model's Out of Sample period starts on 1/31/2022.

The V-Score utilizes Vector Model inputs and outputs. It's training includes 250 model dates randomly selected from 6/1/2005 to 1/31/2021. Forward performance from each model date up to 1 year forward takes the full V-Score training data term up through 1/31/2022. The training period for many strategies discussed in this study begins on 2/1/2022.

The following tickers comprise the complete Vector Model and V-Score coverage universe from which the tickers discussed in this report were drawn:

AA, AAP, AAPL, ABBV, ACGL, ADBE, AMAT, AMC, AMD, AMGN, AMZN, AVGO, AZN, AZO, BA, BAC, BALL, BBY, BHC, BHP, BIIB, BMY, BUD, BXP, CAH, CCL, CDNS, CHTR, CITI, CLF, CMA, CMCSA, CMG, CNC, COST, CPRT, CSCO, CSTM, CTLT, CVS, CYH, CZR, DHI, ELAN, EMB, ETRN, EXPE, FCX, FIS, FITB, FRA, FRCB, FSUGY, GBTC, GE, GILD, GLD, GME, GNRC, GOLD, GOOGL, GS, GSK, GT, GWW, HCA, HD, HLT, HON, HSBC, HYG, IEP, INTC, INTU, IRM, ISRG, JAZZ, JPM, KALU, KEY, KHC, LEN,LLY, LNC, LQD, LUMN, LVS, LW, META, MNST, MOS, MRK, MS, MSFT, MSI, MSTR, MU, MUB, NAVI, NEM, NFLX, NVDA, NVS, NWL, ON, ORCL, ORLY, OXY, PCG, PEP, PHM, POST, PRGO, PWR, QCOM, QQQ, RIO, SBNY, SBUX, SIVBQ, SLV, SNY, SPY, T,TDG, TEVA, TFC, THC, TLT, TMUS, TRGP, TSLA, TXN, UAA, UNH, USB, VCSH, VFC, VICI, VNO, VST, VZ, WDC, WFC, WRK, WYNN, X, XOM, ZION, ZTS

Despite utilization of sklearn's LedoitWolf conditioning, consistent covariance matrix stability for strategies using Trailing for their correlation input required us to take a subset of the 149 tickers in VecViz's coverage listed above when conducting this study. We selected 99 tickers, taking the top, bottom and middle 33 performers for the 4/30/24 thru 9/30/2025 test period that were also present for the entire training period preceding it. The 99 tickers selected are:

AA, AAPL, ACGL, ADBE, AMD, AMGN, AVGO, AZN, AZO, B, BBY, BHC, BXP, CCL, CDNS, CHTR, CLF, CMA, CMCSA, CMG, CNC, COST, CSCO, CSTM, CVS, CYH, CZR, ELAN, EMB, FIS, FRA, FSUGY, GBTC, GE, GILD, GLD, GME, GOOGL, GS, GT, GWW, HCA, HD, HLT, HSBC, HYG, IEP, INTC, INTU, IRM, ISRG, JAZZ, JPM, KALU, KEY, KHC, LLY, LQD, LUMN, META, MNST, MOS, MRK, MSFT, MSI, MSTR, MUB,

NAVI, NEM, NFLX, NVDA, NWL, ORCL, ORLY, PCG, PHM, POST, PRGO, PWR, QQQ, RIO, SBUX, SLV, SPY, T, TDG, TEVA, TLT, TMUS, TSLA, UAA, UNH, VCSH, VFC, VST,WDC, WFC, WYNN, XOM

The VecViz expected return inputs and associated volatility regime framework were formulated via data exploration of VectorModel and V-Score inputs and outputs for the period February 1, 2022 through April 30, 2024.

VecEvents are events or themes that were highly influential on a ticker. The VecEvents that are the basis of VecViz's VecEvent similarity-based correlation were generated by Gemini 2.5 Pro. Prompting for single company ticker VecEvents was done in April 2025. Prompting for ETF ticker VecEvents was done in July 2025. Any VecEvent with a starting "VecDate" (i.e., the start of it's period of influence upon the associated ticker's price) after 12/31/2023 was excluded from this study.

VecViz's "Fingerprint" similarity-based model of correlation (described further in the Appendix) is based entirely on many of the same Vector Model and V-Score inputs and outputs used in the VecViz expected return composite and related regime definition criteria.

During the 4/30/24 thru 9/30/2025 test period, the S&P 500 ETF, SPY, had a CAGR of 22.6% and an annualized vol of 17.9%. In contrast, the average ticker included in the study had a lower CAGR of 19.3%, and an annualized standard deviation of daily returns of 36.4%. The distribution of returns among tickers was highly skewed, with the median ticker CAGR well below the average, at 14.1%. See the box plot and scatter plot below for detail.

Note that the CAGR and volatility for the average ticker as discussed above is quite different from what you will see for returns and volatility of the 1/n portfolio in the study exhibits. The study exhibits reflect daily rebalancing of the 1/n portfolio, which adds 370bps of annualized return before transaction costs. Note also that all strategies presented are likewise rebalanced to their prior rebalance date target weightings for each day of the period studied.

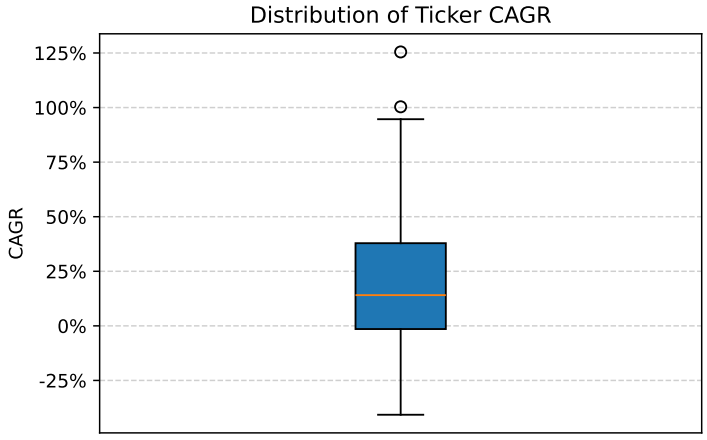

(a)

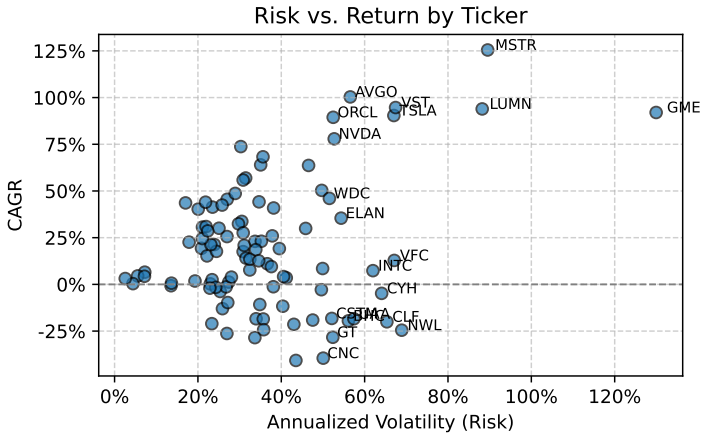

(b)

Though it featured significant market volatility and dispersion in ticker returns, the 17 months spanned by the test period isn't a great deal of time. For that reason, the Appendix includes, a repeat implementation of the full

study for the 3/31/2023 through 9/30/2025 period. Note that this 30 month period includes 13 months of that were used to specify its VecViz's return composite indicator and its VecEvent and Fingerprint similarity-based correlation metrics. Please also note that the period between the 1/31/2022 end of the V-Score training period and the 3/31/2023 start of this supplementary study provided in the Appendix was needed to meet the minimum lookback window requirements of the VecViz regime-based expected return composite selection framework and all the Trailing based metrics.

Input Detail

Classical MVO relies on three core inputs: expected returns, risk (volatility/covariance), and diversification (correlation structure). An objective (e.g., maximize expected return) is required, and constraints (e.g., max individual exposure weight, max expected risk) are typically included.

Expected Return

The objective of each MVO strategy considered here is to maximize expected return, where expected return is defined as either Trailing 252d returns or VecViz's volatility regime based composite return indicator.

Volatility regime based investment tactics have been a subject of much academic and industry research. VecViz's volatility regime definition is bit uncommon, based upon violation of expected volatility thresholds, in the form of 95% VaR (the 95th percentile of Value at Risk) to the downside and 99% OaR (the 99th percentile of Opportunity at Risk) to the upside, breakage rates.

The 95% and 99% threshold percentiles for VaR and OaR, respectively, were determined via exploration of the period between 2/28/2022 and 4/30/2024, and are consistent with the notion that most investors are more readily inclined to experience fear of losses than fear of missing out.

VecViz's regime based composite return indicator is specified via a two step process. The first step is identifying the regime. VecViz characterizes volatility regimes into hi/mid/lo buckets for both 95% VaR and 99% OaR breakage, resulting in a total of 9 regimes.

The second step is identifying the VecViz indicator most closely associated with higher returns for the regime identified. All expected return composites utilize a standardized V-Score percentile, with most adding several standardized Vector Model input or output feature percentiles to it. The VecViz expected return composites considered, and composite usage by and across regimes are discussed further in the Appendix.

Note: Expected return composite performance was evaluated on the basis of next day performance at the date of original publication (10/22/2025). Starting on 11/11/2025 we started evaluating composite performance on the basis of 10 day forward performance and republished the back test to reflect that methodology change. The back test utilized dates no later than 11 days prior to each rebalance date to select expected return composites, to avoid look-ahead bias.

Constraints

Each MVO strategy is subjected to the same set of constraints. These include:

1) maximum expected volatility
2) maximum weight per ticker
3) the requirement that all ticker weights are > 0
4) the sum of all ticker weights is $<= 100\%$.

Covariance

In MVO, portfolio risk is a function of the covariance matrix of ticker returns. Two of the MVO strategies in this study, S1 and S4, rely upon a covariance matrix based entirely on trailing 252d returns. The other strategies assemble the covariance matrix via the product of ticker standard deviation and pairwise ticker correlation (cov(x,y) = stdev(x)*stdev(y)*corr(x,y)), employing one of two inputs for the former (Trailing 252d and VecViz 99D), and one of three inputs for the latter (Trailing 252D, VecViz VecEvent Similarity, and VecViz Fingerprint Similarity).

Ticker level volatility

The inputs for ticker standard deviation include trailing 252d returns and VecViz's "99D_Ret", which we divide by 2.326 in an attempt to scale it at least somewhat similarly to the alternative.

Pairwise Ticker Correlation

The inputs for ticker correlation include trailing 252d returns and two distinct VecViz based correlation solutions sets: VecEvent similarity and "fingerprint" similarity.

VecEvent similarity applies a variety of techniques, including sklearn's TF-IDF (Term Frequency-Inverse Document Frequency) vectorizer, K-Means clustering, and cosine similarity to derive textual benchmarks from a set of ticker VecEvents and to calculate each ticker's exposure to each benchmark. VecEvent correlation is the correlation of each ticker's exposure to these benchmarks. Calibration of the number of textual benchmarks to utilize was performed by minimizing the mean adjusted error to Pearson correlation for corresponding ticker pairs for the period February 1, 2022 through April 30, 2024. The resulting pairwise VecEvent based correlations were held constant through the 4/30/2025 through 9/30/2025 test period. The process of calculating VecEvent similarity is described further in the Appendix.

"Fingerprint" refers to a ticker's Vector Model and V-Score inputs and outputs, which together comprise its "fingerprint", from a VecViz perspective. We consider 21 such metrics, each standardized to a max/min scale for

the day prior, for each ticker on each rebalance date, and take the correlation between them across ticker pairs. The process of calculating "Fingerprint" similarity is described further in the Appendix.

Calculation Detail

Covariance Conditioning (or not)

We use skLearn's LedoitWolf function to build the covariance matrix in a robust manner when relying entirely on trailing 252d returns for correlation and ticker level volatility inputs.

When utilizing VecViz inputs for the ticker level volatility or the pairwise ticker correlation as components of the covariance matrix, LedoitWolf is not used.

Optimization engines utilized

We then feed the covariance matrix along with the constraints and the objective into a loop that attempts to find the optimal solution by applying a series of optimizers, starting with cvxy's cp.ECOS, then, if that fails attempting cp.OSQP, and then, if that fails, cp.SCS. If no solution is found then all tickers are weighted zero.

Point in Time Related Methodological Detail

This study evaluates each strategy using three rebalance date frequencies: 10d, 21d, 63d. If day "d" is a rebalance date then the MVO is ran on day d using information for the period up to d-1. The weights generated by the MVO are ascribed to the closing price on day d and are reflected in performance for day d+1 onward, until the next rebalance date.

Note that our calculations do not capture "drift" between rebalance dates. Instead, we apply the same optimized weighting for each day until the next scheduled rebalance date occurs, at which point the optimization is reran. Therefore, we effectively rebalance ticker weightings for all strategies every day, to either a newly refreshed optimization solution set (reflecting information known as of the prior day's close) or the most recently furnished version thereof.

VecViz inputs for expected forward returns are volatility "regime" driven. This entails examination of a rolling history of regime classification, metric levels, and forward performance. The expected return metric best suited for d-1's regime is determined by examining forward 1 day performance for all composite return metrics during d-1's regime for all identified regimes up through day d-2.

Performance Metrics

SummaryZ Score

The SummaryZ Score is intended to facilitate easy comparison between MVO strategies, and more importantly, between the metrics contributing to those strategies, across a number of important performance dimensions while maintaining a neutral bias with regard to risk aversion vs return seeking investor proclivity.

More specifically, the SummaryZ score is an equally weighted sum of "Z scores" calculated for each of the following six metrics across all portfolio scenarios (except SPY and 1/n) defined as follows:

```
Z_Sum = Z_PR + Z_MDD + Z_SR + Z_CR + Z_A + Z_KPV
```

where:

1) PR: Average annualized portfolio price return
2) MDD: Max draw down of cumulative portfolio price return
3) SR: Sharpe Ratio (PR / Standard Deviation of Price Returns)
4) CR: Calmar Ratio (PR / MDD)
5) A: Alpha of portfolio price return to SPY, MTUM, VLUE (the S&P 500 ETF, the Momentum Factor ETF, the Value Factor ETF)
6) KPV: Kupiec Test Statistic P-Value, which reflects the probability that the portfolio's 99% VaR, as implied by its volatility constraint, (assuming normality, and independent, identically distributed daily returns) was well specified.

Taken together, we see that these components produce a SummaryZ score is 1/3rd return oriented (PR and A), 1/3rd risk oriented (MDD and KPV), and one third risk/return ratio oriented. At a high level one could say that it is 50% risk oriented, 50% return oriented.

Alternatively, one could also notice the breadth of perspective with which both the issue of risk and return is addressed. With regard to return, it considers not just magnitude (explicitly and via the numerator of the Sharpe and Calmar ratios) but also the independence of the return generated from factors (via Multi-Factor Alpha). With regard to risk it considers not just volatility (the denominator of the Sharpe ratio), or maximum draw down (explicitly and within the Calmar Ratio) but also how consistent the volatility experienced was with model expectations (Kupiec P-Value).

We do not consider transaction cost in SummaryZ nor penalize for turnover, but we do disclose turnover incurred at rebalance dates in the Appendix. Note that this turnover measure excludes the rebalancing that occurs to maintain ticker weights at the most rebalance date specified levels between rebalance dates. Turnover for all strategies, including 1/n, is therefore understated.

Multi-Factor Alpha

The Multi-Factor Alpha regresses the intercept of the daily return of each strategy against the corresponding returns of MTUM, VLUE, and SPY, multiplied by 252. These tickers represent the iShares MSCI USA Momentum Factor ETF, the iShares MSCI USA Value Factor ETF, and the SPDR S&P 500 Trust ETF, respectively.

Calmar Ratio

The Calmar Ratio is the ratio of Average Return to Max Draw Down.

Kupiec Test Statistic P-Value

To be clear, we are not measuring the portfolio's VaR and testing it via the Kupiec statistic. Instead, we are using Kupiec to test adherence of the portfolio generated by each strategy to the volatility constraint, relying on the assumption that returns will be normally, indepenedently and identically distributed. Thus, for strategy iterations involving a 20% annualized volatiltiy presumption, we consider 99% VaR Breakage to occur when portfolio daily return is less than (-20%/Sqrt(252))*2.326 = -2.93%. The closer the breakage rate for a strategy is to 1.00% the higher the Kupiec Statistic P-Value will be. Thus, overly conservative portfolios with 0% breakage are penalized as much as portfolios with 2.00% breakage. Ultimately, it measures how well the portfolio risk estimate was calibrated.

Important considerations:

1) Past performance is no guarantee of future results. None of the content in this report is investment advice or an offer to buy or sell securities. VecViz is not an SEC investment advisor or broker-dealer. The staff of VecViz actively transacts in securities tied to many of the tickers discussed in this report. See VecViz's Terms and Conditions for more context and detail at https://vecviz.com/termsand- conditions/
2) We are not considering transaction costs in the SummaryZ metric, though we detail turnover (on a 2 way, round trip basis) in the Appendix.
3) The time period of May 1, 2024 – September 30, 2025 is brief for this sort of study. However, it did include a high level of volatility and return dispersion as detailed in the "Data" section earlier. For additional perspective, skewed though it may be, see the Appendix for results that include the Training Period and the Test Period, spanning 4/1/2023 through 9/30/2025.
4) There are a number of techniques which are known to enhance the performance of trailing return based volatility analytics such as exponential time decay of observations that we did not implement in this study. The relative performance of VecViz analytics would not have been as strong as it is depicted here had we utilized those techniques.

5) We should explore how different lookback windows, rebalance dates, and ticker sets would affect the results.

Motivation and Acknowledgements

VecViz's Vector Model is a model of ticker level price probability. From it we produce a suite of analytics - VaR, OaR, Option Fair Value Estimates, etc.. Though the design of the Vector Model emphasizes broad cognitive accessibility via linkage to visual calculation inputs (the Vector Set channels) and narrative (the VecEvents linked to those channels) as much as accuracy, we report exhaustively each month on the accuracy of the metrics produced. The "Reports" section of vecviz.com features a separate report for each metric, and each report has a voluminous appendix that allow the reader to interrogate the summary results.

Despite all the effort invested into assembling the report for each of these metrics, we readily acknowledge that the context of portfolio optimization subject to a range of realistic constraints produces a more informative, meaningful, and of course, concise, summary of our metric performance. Doing so has long been on our "to do" list, but remained out of reach until this summer due to limited resource availability.

The extension of the Vector Model's ticker level analytics into Portfolio Analytics was made possible by the contributions of a few interns from Cornell's School of Financial Engineering Mohamed Azahriou, Guangyou Zhou, Yushang Wu, and a recent graduate of Lehigh's Financial Engineering program, Fathmat Samira Bakayoko. They were instrumental in VecViz feature exploration, the expected return regime framework development, the development of both our VecEvent and fingerprint similarity-based correlation metrics, and the overall framework for running the 300+ grid search scenarios presented in this report and meticulously aggregating and reporting the results. We thank them for their contribution.

Results

Objective 1: Comparing the value add of VecViz and Trailing MVO inputs

Performance of the strategies that each VecViz and Trailing input contributes to, across all 27 constraint and rebalance frequency combinations, is summarized in the table entitled "Average Performance Metrics by MVO Input Variable", that follows in the pages ahead.

VecViz's "99D_Ret" input for ticker volatiltiy is the most valuable MVO input considered in this study, with an average SummaryZ score of +2.10. Following 99D_Ret is Trailing's expected return input with a score of +1.45 and VecViz's VecEvent-similarity based correlation, with a score of +1.25.

Driving 99D_Ret's performance during the study period was its tendency to have smaller drawdowns and VaR breakage more consistent with vol constraint levels. That said, the Annualized Return and Alpha outcomes of the strategies it was an input to were both among the top 3 of the 7 inputs considered. Those return outcomes in conjunction with the leading volatility profile resulted in the top outcome for Calmar Ratio and a near miss on the top outcome for Sharpe Ratio.

We note here that the supplemental study included in the Appendix, which starts on 3/31/2023, has an identically structured table. It shows that for the 3/31/2023 through 9/30/2025 period VecViz's VecEvent-similarity based correlation was the the most valuable input, with a score of +1.51. Following close behind were Trailing 252d's expected return metric at +1.14 and VecViz's "99D_Ret" metric, at +0.88. Some thoughts on the shift in the relative value add of these metrics relative to the test period:

1) We ascribe the VecEvent-similarity based correlation's jump to the leading SummaryZ position for this expanded period largely to the fact that it was tightly calibrated upon realized Pearson correlation for the 3/1/2022 through 4/30/2024.

2) We ascribe Trailing's expected return input shifting ahead of VecViz's 99D_Ret in SummaryZ terms to the extremely strong price return performance of momentum based strategies for this expanded period. VecViz's 99D_Ret advantage in all risk related metrics over Trailing metrics for this expanded period was similar or greater than it was in the fully out of sample study test period.

Performance Metrics by MVO Input Variable

Average performance metric values are calculated across all strategies that include the given optimization input variable. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.'

Average Performance Metrics by MVO Input Variable
for model dates and 1d fwd perf starting 2024-05-01 and ending 2025-09-29.

Input Variable	SummaryZ	Ann_Ret	MaxDD	Sharpe	Calmar	alpha_ann%	Kup.-Pval
Vol_VV	2.59	23.8%	-15.4%	1.61	1.63	9.8	0.26
Correl (VE)_VV	1.35	24.6%	-19.7%	1.51	1.40	15.4	0.13
Ret_T252d	0.85	25.3%	-21.1%	1.51	1.36	12.4	0.11
Correl (FP)_VV	0.17	22.8%	-19.2%	1.46	1.33	4.2	0.18
Ret_VV	-0.75	20.5%	-18.0%	1.32	1.21	2.3	0.19
Correl_T252d	-1.38	21.1%	-19.9%	1.27	1.14	2.5	0.14
Vol_T252d	-2.49	21.9%	-23.7%	1.22	0.94	5.0	0.04

Objective 2: Comparing the performance of VecViz based portfolios to the "1/n" portfolio and SPY

You can compare the performance of all 12 strategies against each other and SPY and the "1/n" portfolio across all constraint levels and rebalancing frequencies in the table that follows.

Performance relative to the 1/n portfolio:

The performance of the "1/n" portfolio and the performance of each strategy studied reflect transaction cost free daily rebalancing, making them fairly comparable.

1) The "100% VecViz" strategies, S10 and S14, both solidly beat 1/n on a SummaryZ score basis, in both the fully out of sample test period and the extended period presented in the Appendix. In contrast, the 100% Trailing portfolio, S1, underperformed 1/n in the test period, but outperformed it in the extended period presented in the Appendix.
2) Four of the eleven strategies containing VecViz inputs (S4, S9, S11, and S13) underperformed 1/n during the test period. The common trait between them was that they relied upon Trailng 252d ticker volatiltiy instead of VecViz's "99D_Ret" metric. In the extended period presented in the Appendix just two of the eleven strategies containing VecViz inputs (S4 and S9) underperformed 1/n.

Performance relative to SPY:

Some important considerations to keep in mind when considering Strategy performance vs. SPY:

a) SPY is itself one of the factor's in the Multi-Factor alpha calculation, which causes it to have zero alpha. Since most strategies have positive Multi-Factor alpha as measured we attempt to mitigate this issue by ascribing SPY a "0" Multi-Factor Z-score.

b) Differences in constituents - only sixty of the nintey nine tickers included in our study's ticker universe are presently constituent tickers of the SPY ETF. This issue is somewhat mitigated by the fact that only 39 of the 99 tickers studied had a CAGR that outperformed the SPY's (during the test period).

c) Unlike any of the strategies studied, the SPY's return reflects the cost of its rebalancing. This issue is mitigated by the fact that the S&P 500 index related turnover is very low, typically in the range of 2-5% per year, and SPY typically has a very low tracking error to the S&P 500 Index, despite share redemption and creation related turnover activity assocated with being an ETF.

All that said, here are the key SPY comparison results:

1) The "100% VecViz" strategies, S10 and S14, both solidly beat SPY on a SummaryZ score basis during the fully out of sample test period and the extened period presented in the Appendix. In contrast, the 100% Trailing portfolio, S1, underperformed SPY in the test period, but outperformed it in the extended period presented in the Appendix.

2) In all, nine of the eleven strategies containing VecViz metrics outperformed SPY on a SummaryZ basis, for the fully out of sample test period and the extended period featured in the Appendix.

Additional Strategy Performance Observations:

The table shows that S8 was by far the strongest strategy. S8 utlizes Trailing 252d for its return input and VecViz for its volatility and correlation inputs. More specifically, for correlation it utilized VecEvent similarity based correlation. Given the ambiguity about "point-in-time" when dealing with LLM output about past events, we note that S12 was a close second to S8 in terms of consistently strong SummaryZ results. S12 also utilized Trailing 252d for its return input and VecViz for its volatility inputs, and VecViz's Fingerprint similarity for its correlation input.

"SummaryZ" Performance Score Table

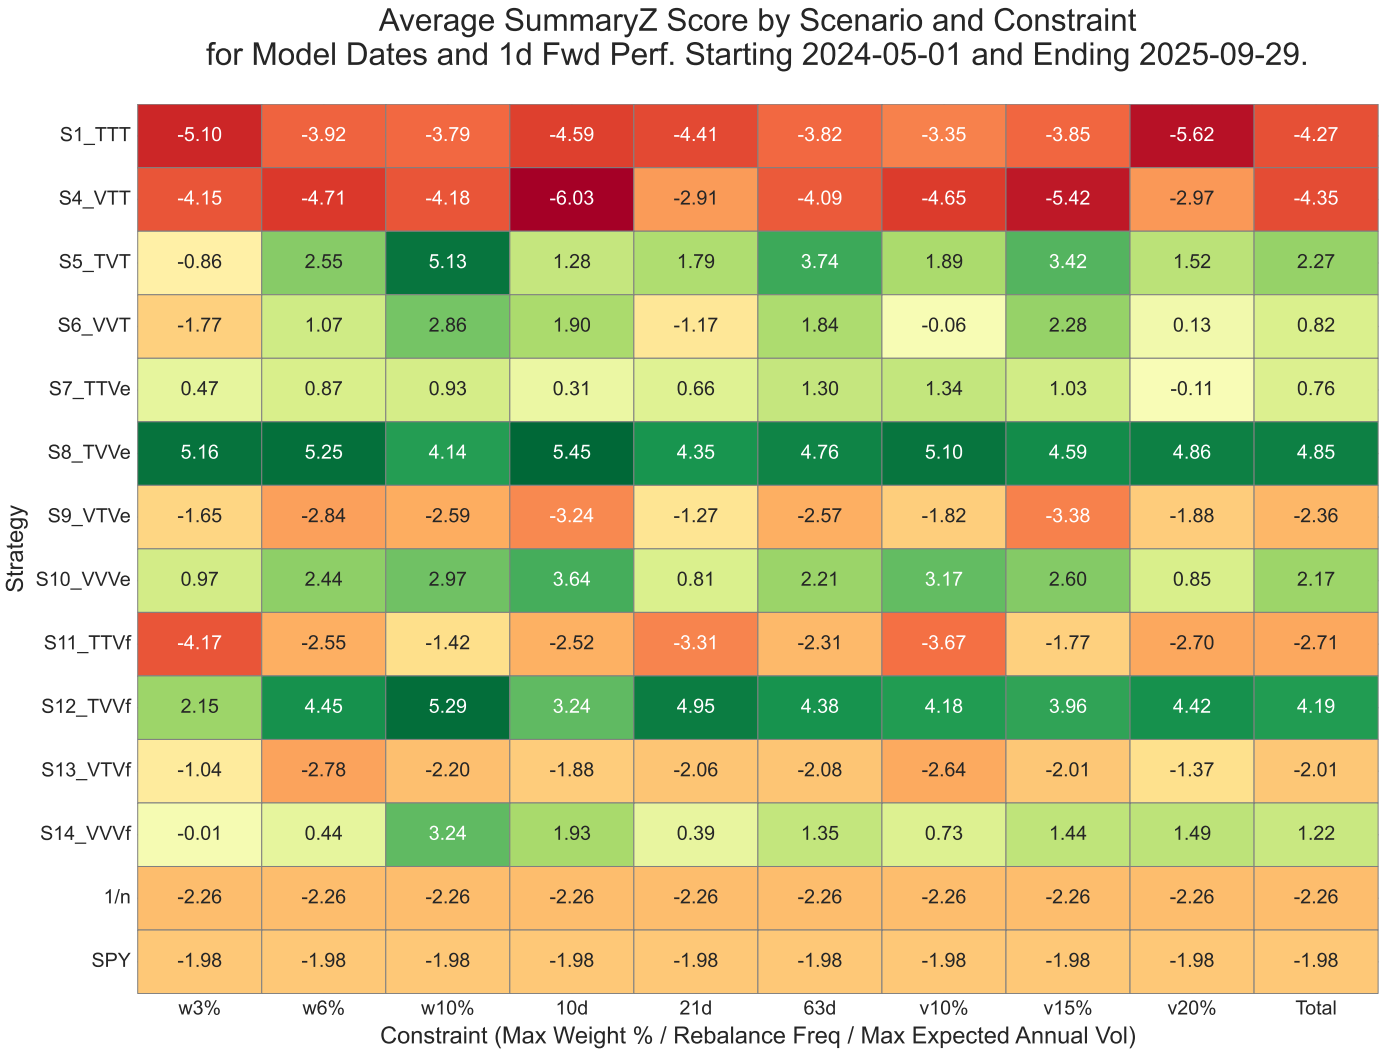

SUMMARYZ score, a sum of Z-scores calculated across all strategies and constraint combinations for six metrics: Annualized Return, Sharpe Ratio, Max Drawdown, Calmar Ratio, 99% VaR Kupiec P-Value, and Multi-Factor Alpha. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: May 1, 2024 through Septmeber 30, 2025

Appendix: Strategy - Constraint Scenario Cumulative Realized Return Paths

```
Generating plot...
```

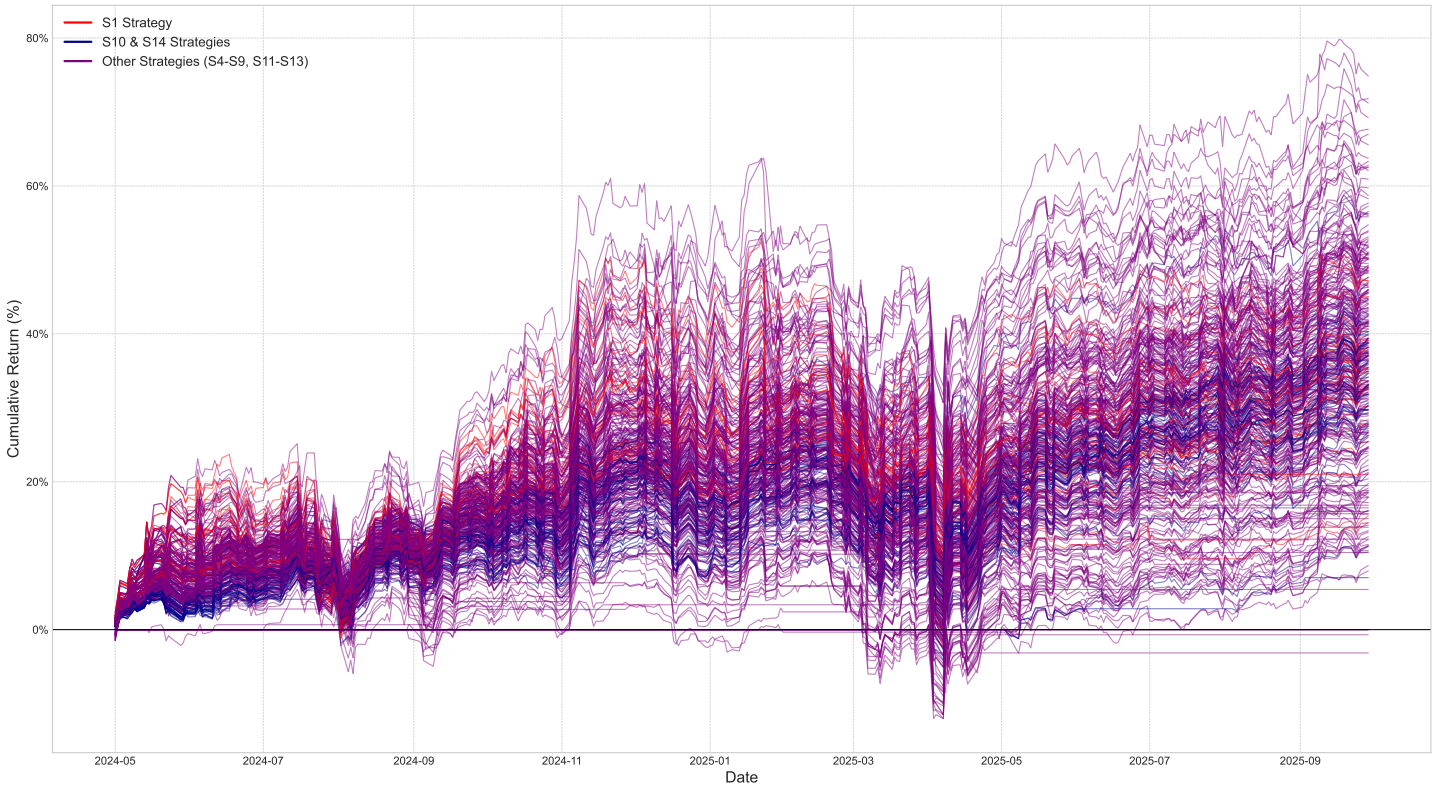

Figure 2: Cumulative Returns of All Grid Search Strategies

Appendix: Realized Return vs. Risk Scatter Charts for 5/1/2024 thru 9/30/2025

Squares= VecViz Return, Circles=Trailing 252d Return | Red=All Trailing, Purple=Hybrid, Blue=All VecViz

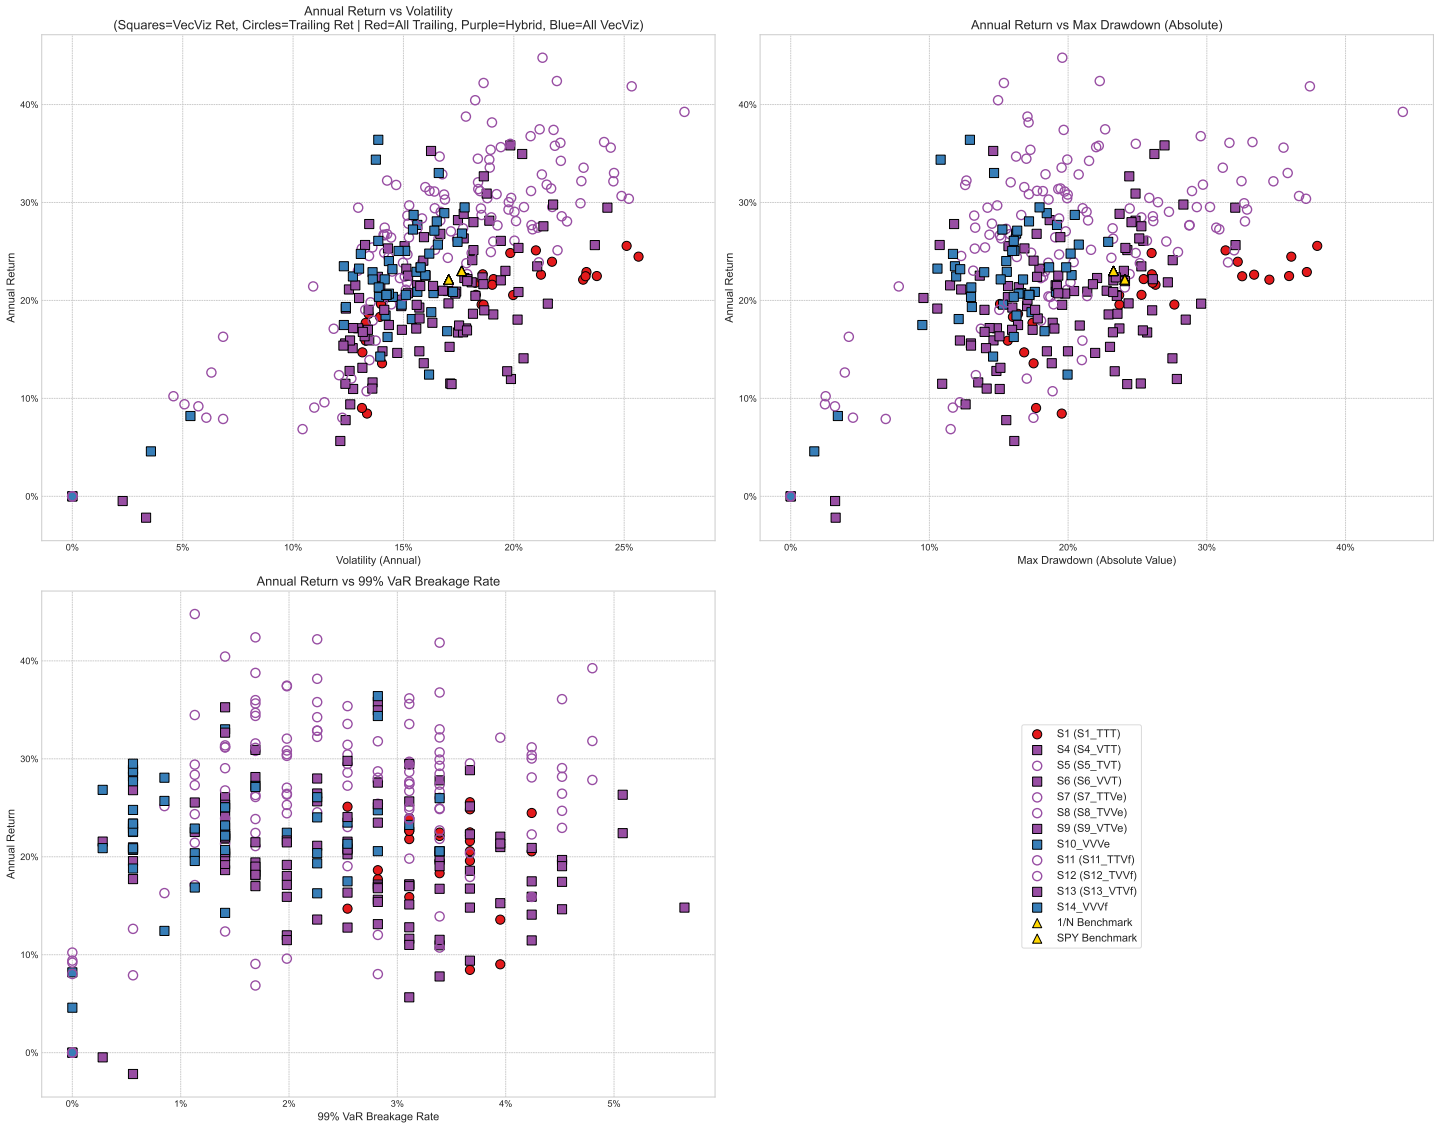

Appendix: Annualized Return (included in SummaryZ)

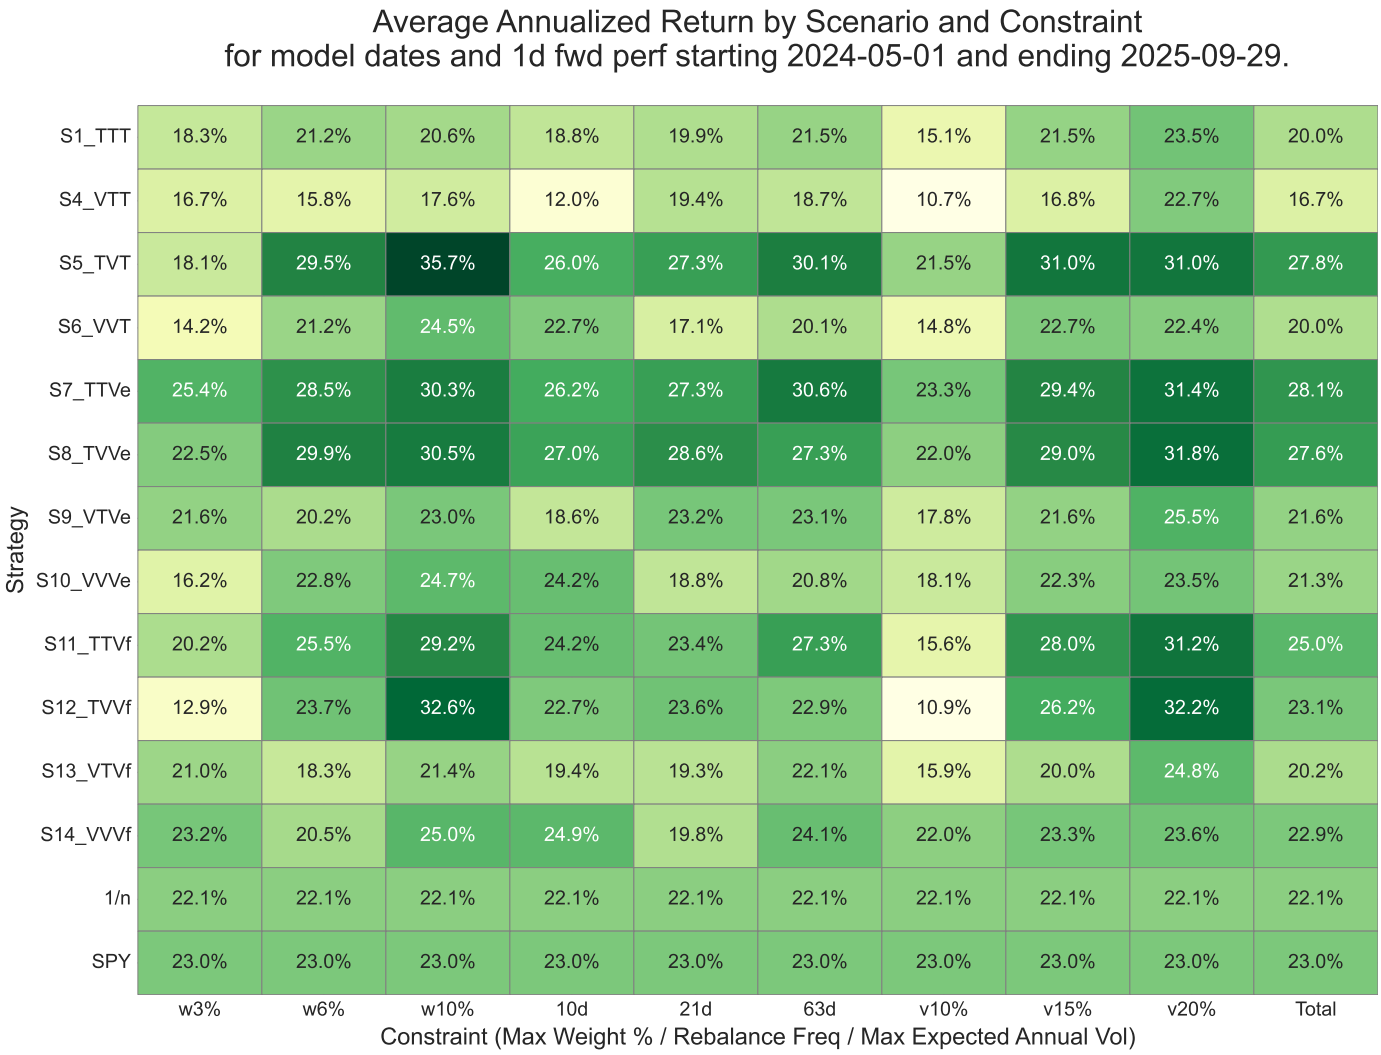

Ann_Return, is the annualized average daily strategy return. Returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Sharpe Ratio (includded in SummaryZ)

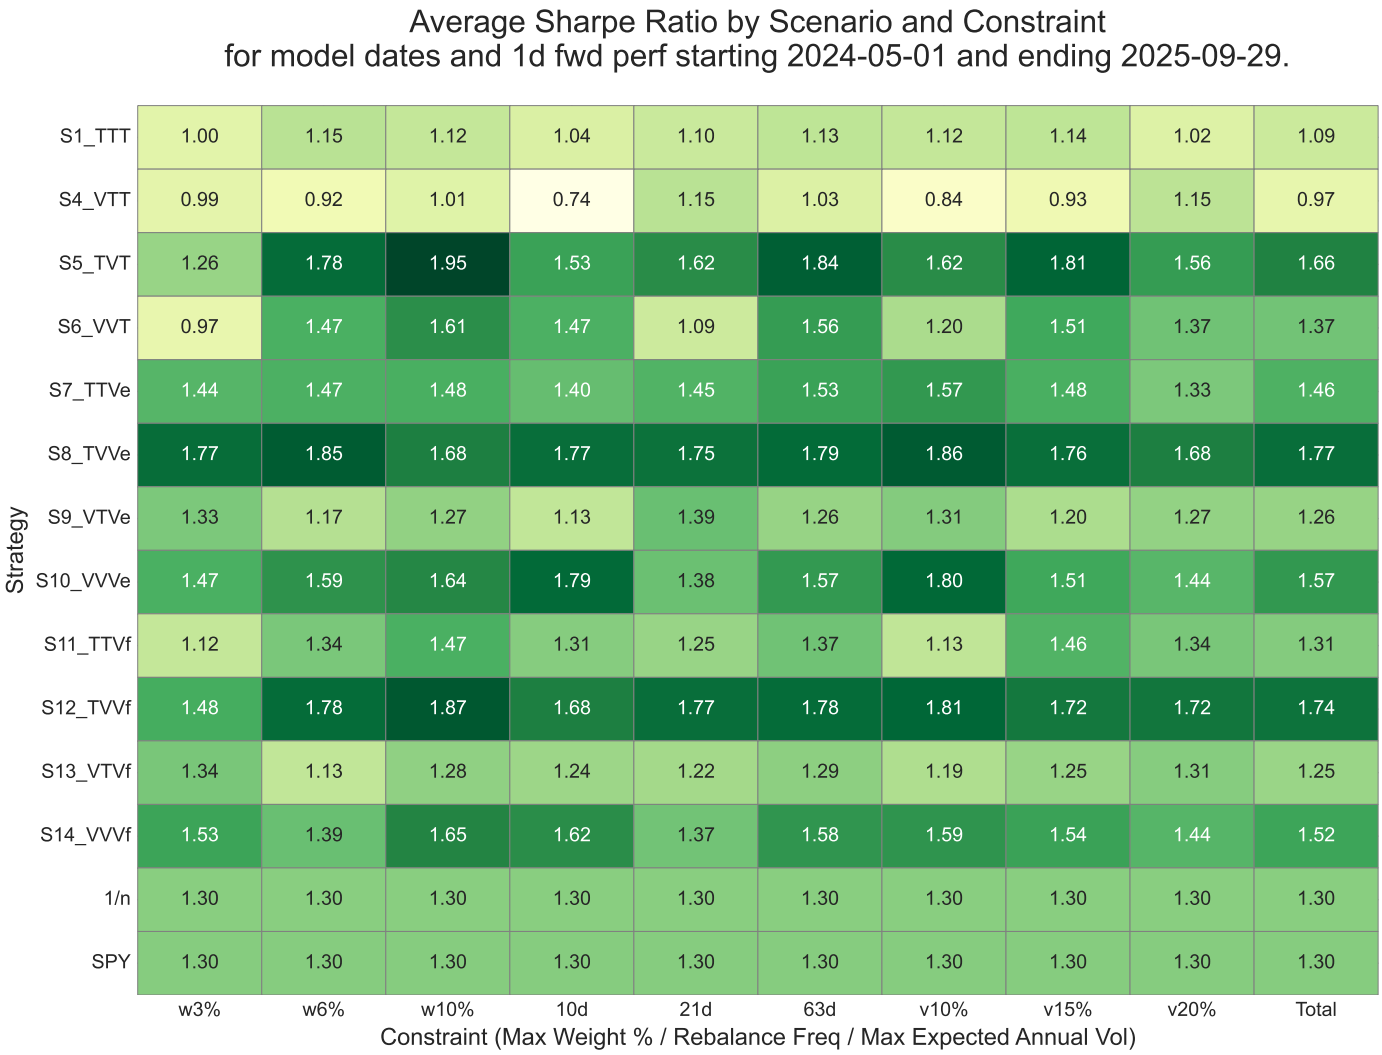

Sharpe Ratio is the annualized daily average strategy return / standard deviation of daily strategy returns. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Max Draw Down (included in SummaryZ)

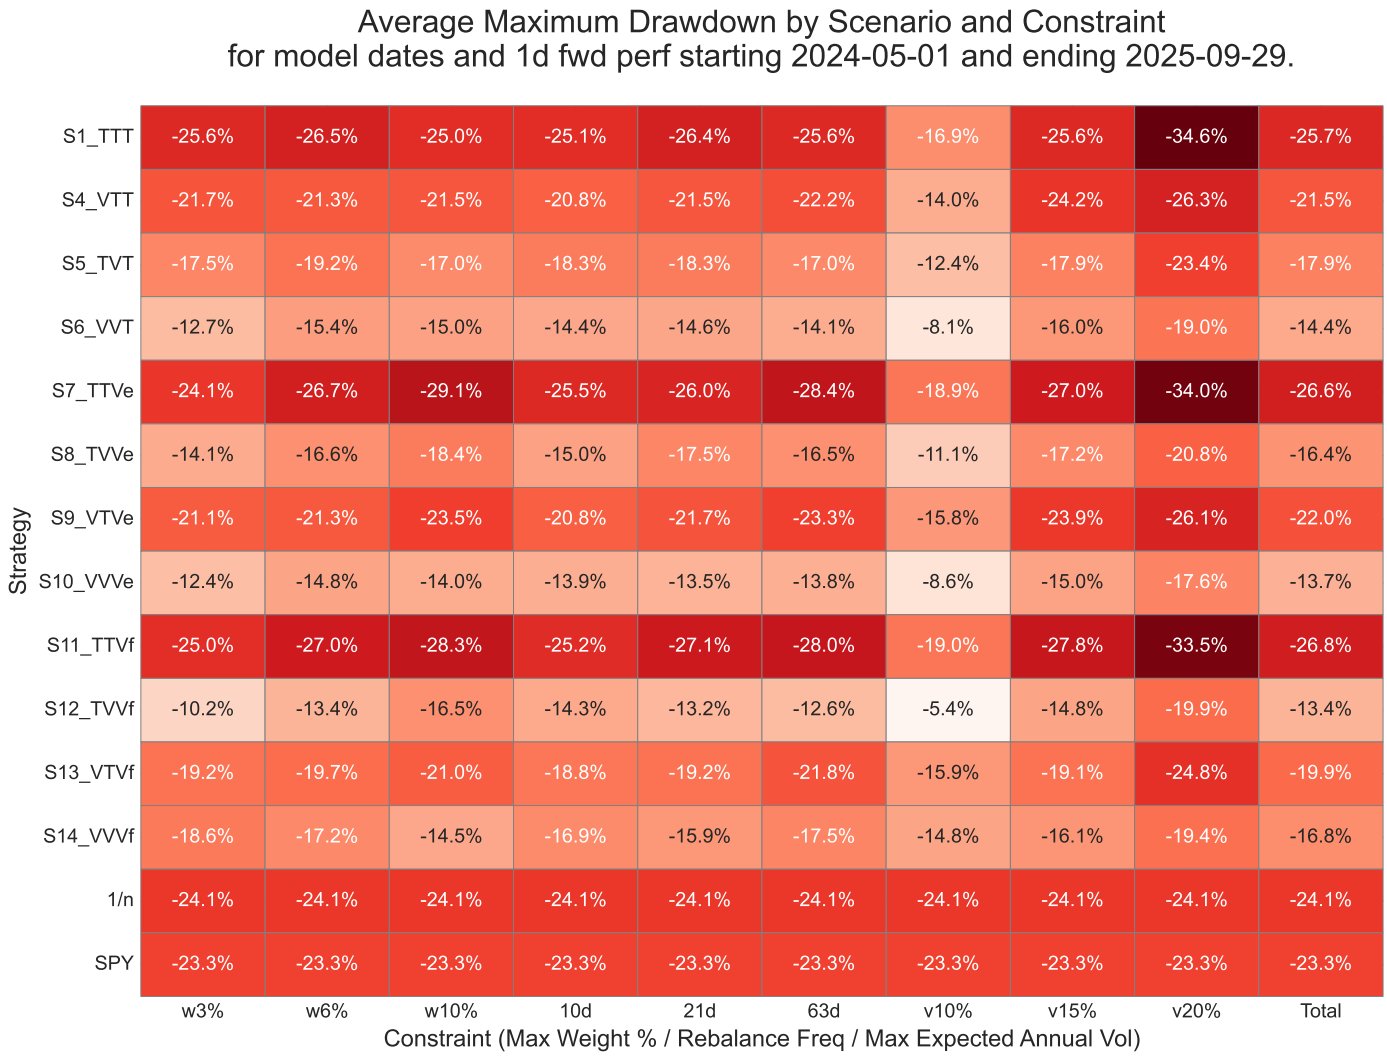

Max Draw Down (MaxDD), is the greatest decline from a peak in cumulative strategy return to a nadir in cumulative strategy return. Returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Calmar Ratio (included in SummaryZ)

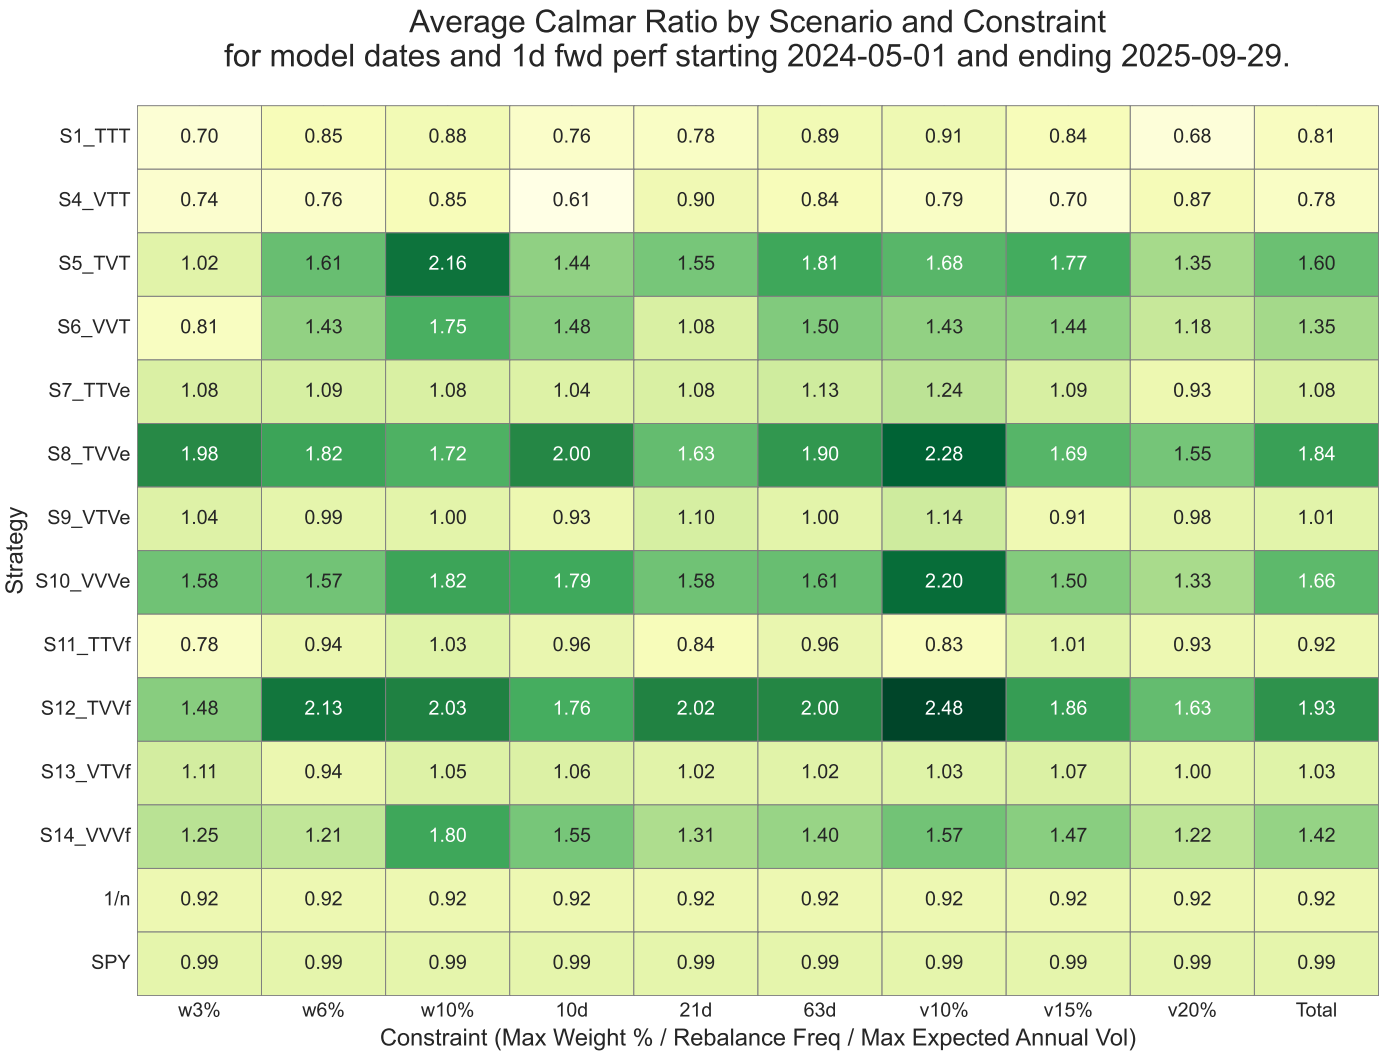

Calmar Ratio is the annualized daily average strategy return / Max Drawdown of cumulative daily strategy returns. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Multi-Factor Annualized Alpha (included in SummaryZ)

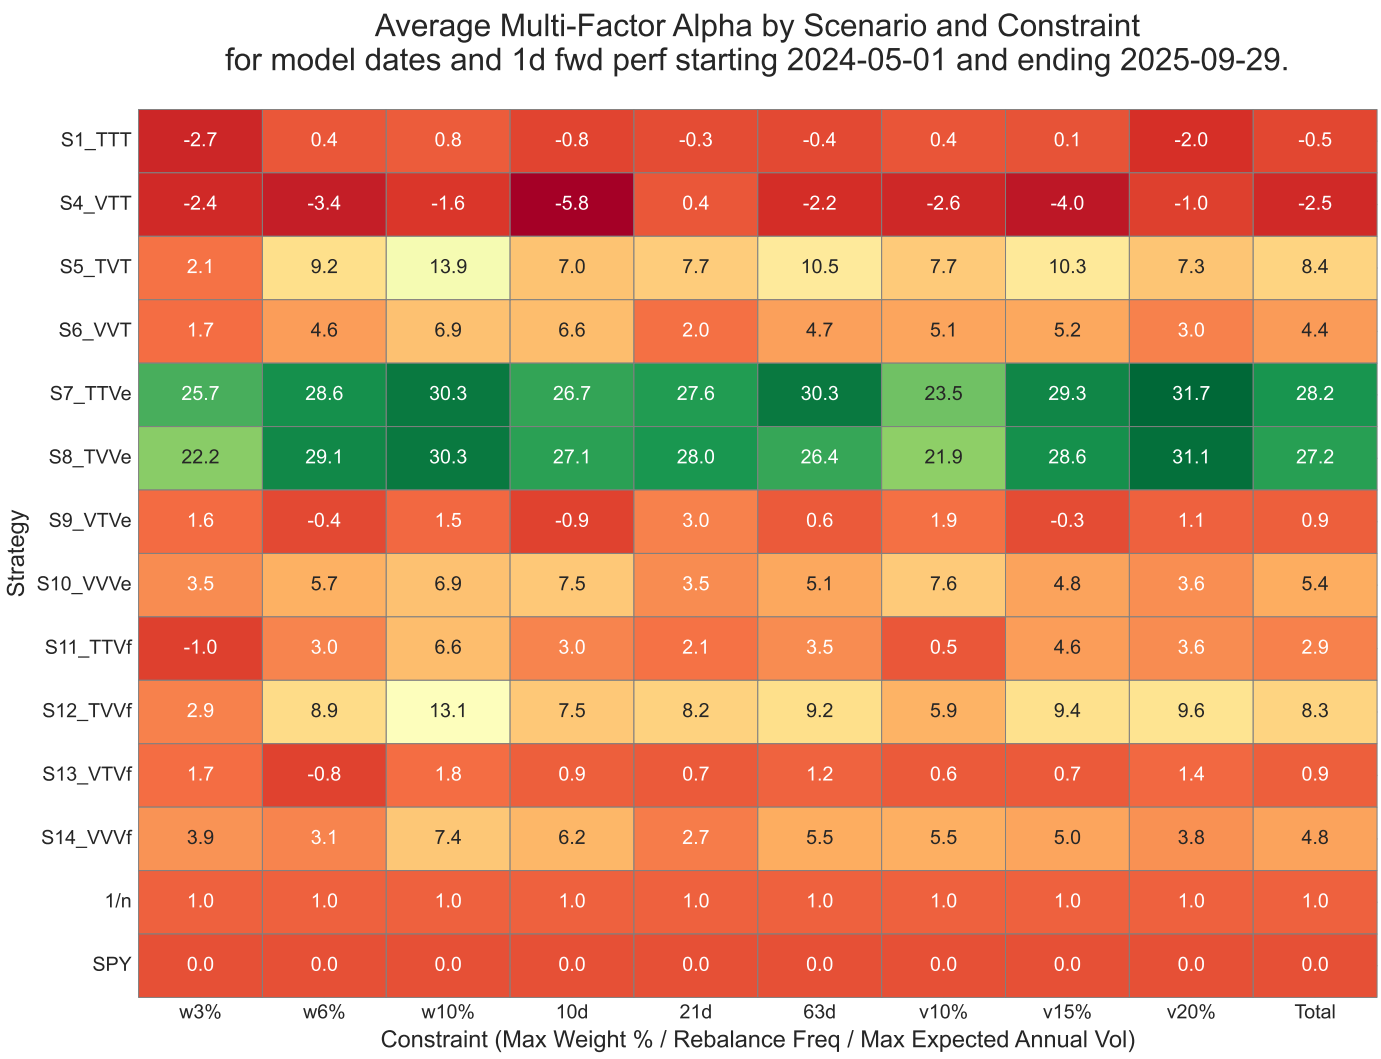

Multi-Factor Alpha is the annualized value of the intercept of daily strategy returns regressed upon the corresponding daily returns of SPY, MTUM, VLUE. Returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: 99% VaR Kupiec Test P-Value (included in SummaryZ)

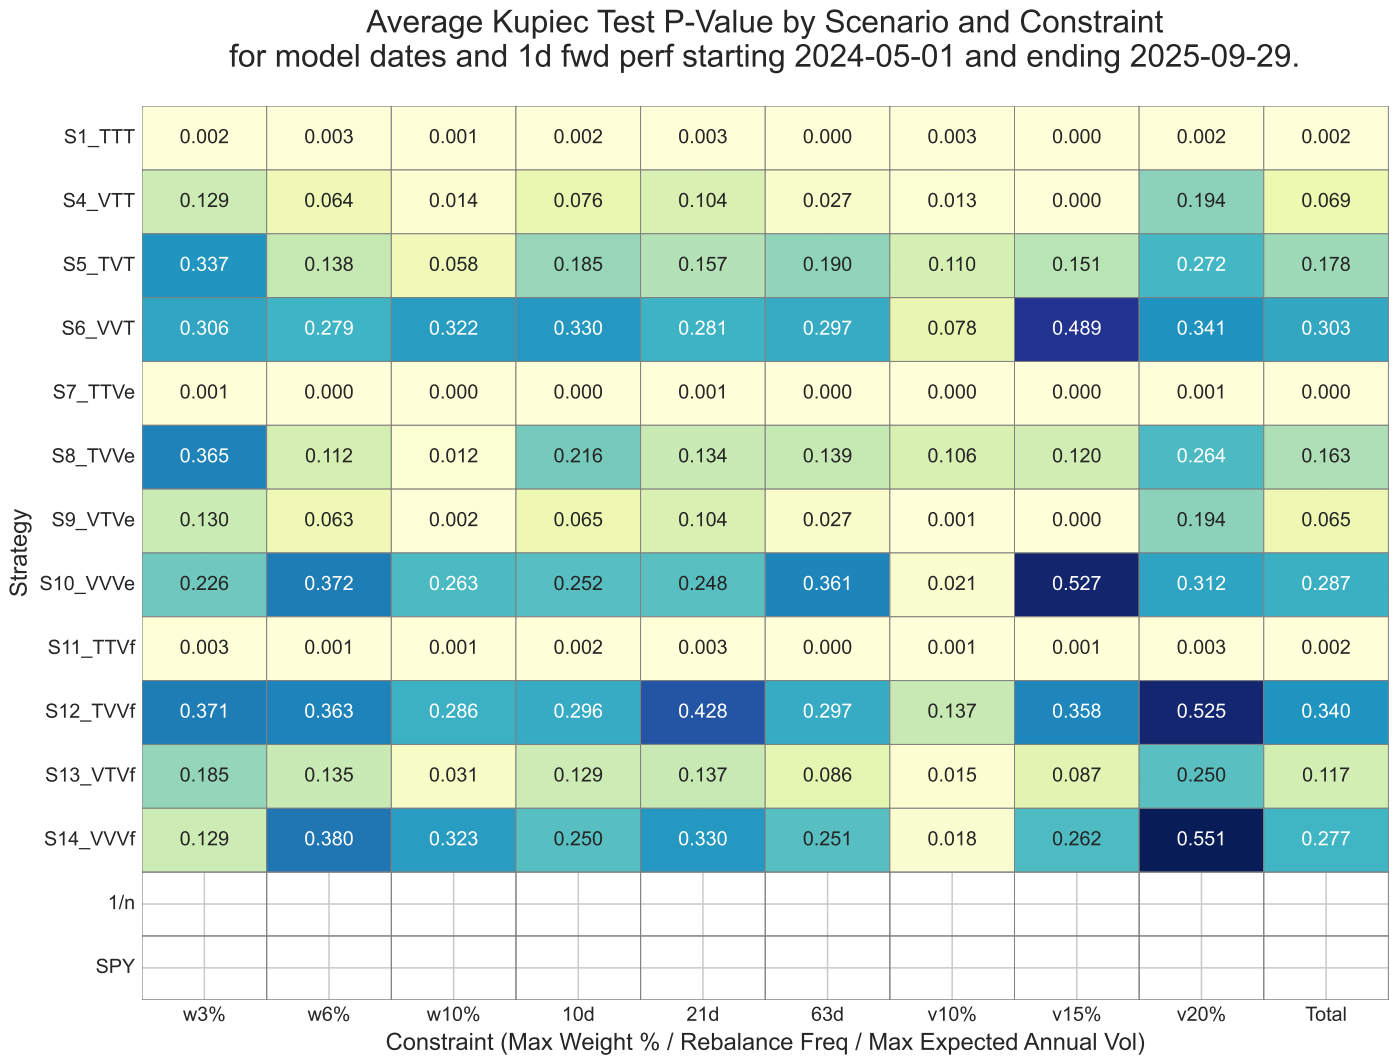

The P-Value of the Kupiec Proportion of Failures test statistic represents the probability that the null hypothesis that the 99% VaR breakage rate is consistent with expectations is true (higher Kupiec P-Value is better). ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Annualized Volatility

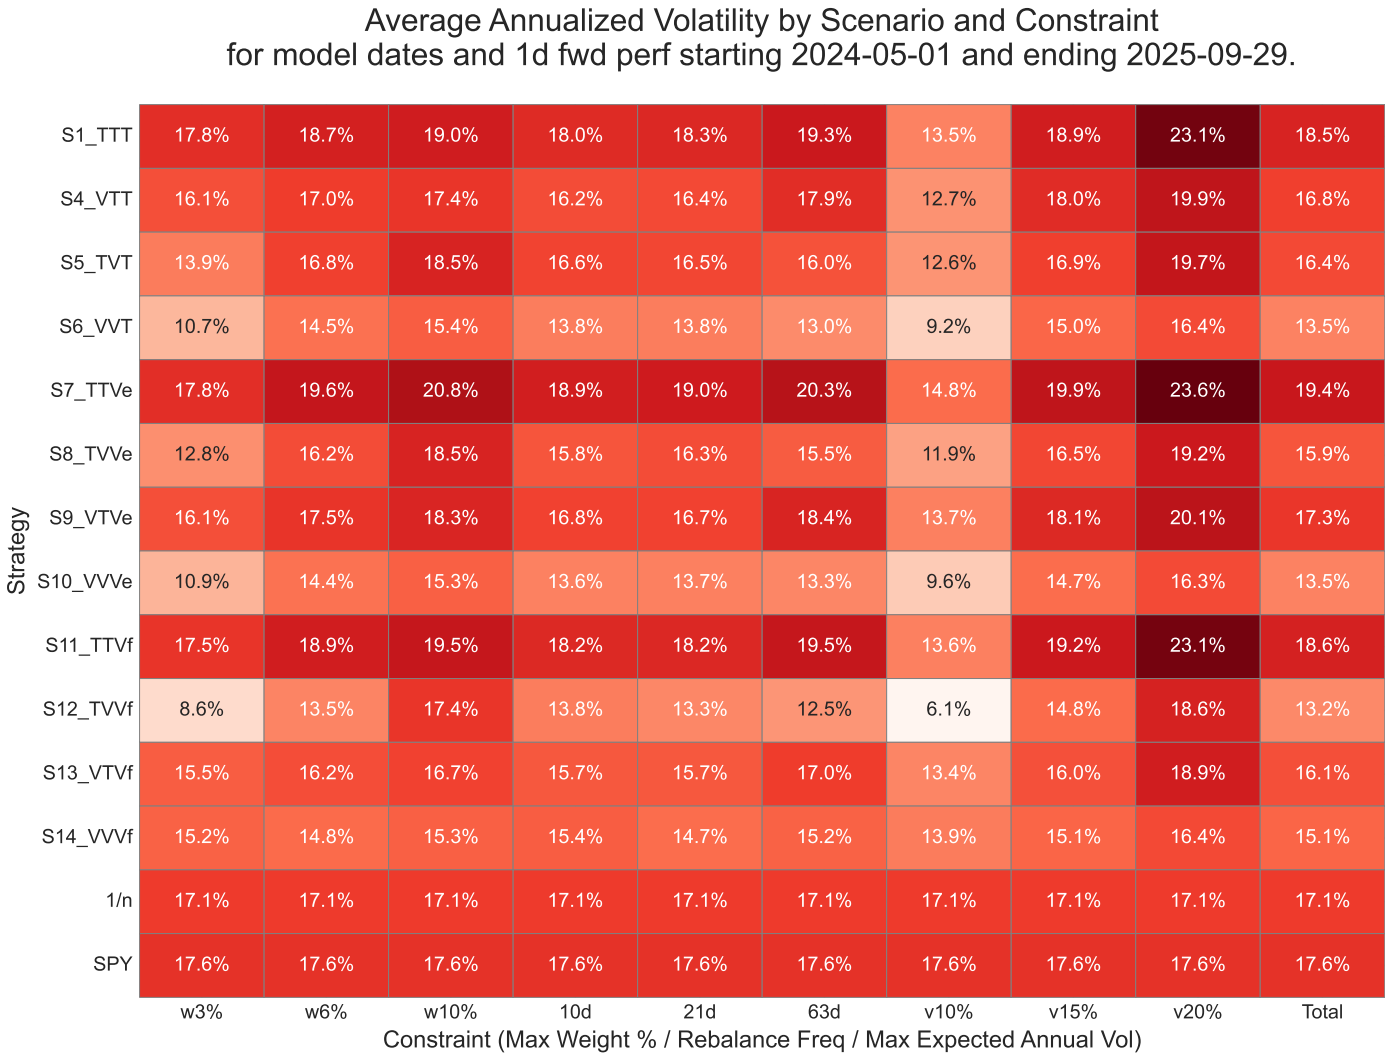

Vol, is the annualized standard deviation of daily strategy return. Returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. Annualized bolatility is Not included in SummaryZ because it is well represented in Sharpe Ratio, Max Draw Down and Kupiec P-Value. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: 99% VaR Breakage

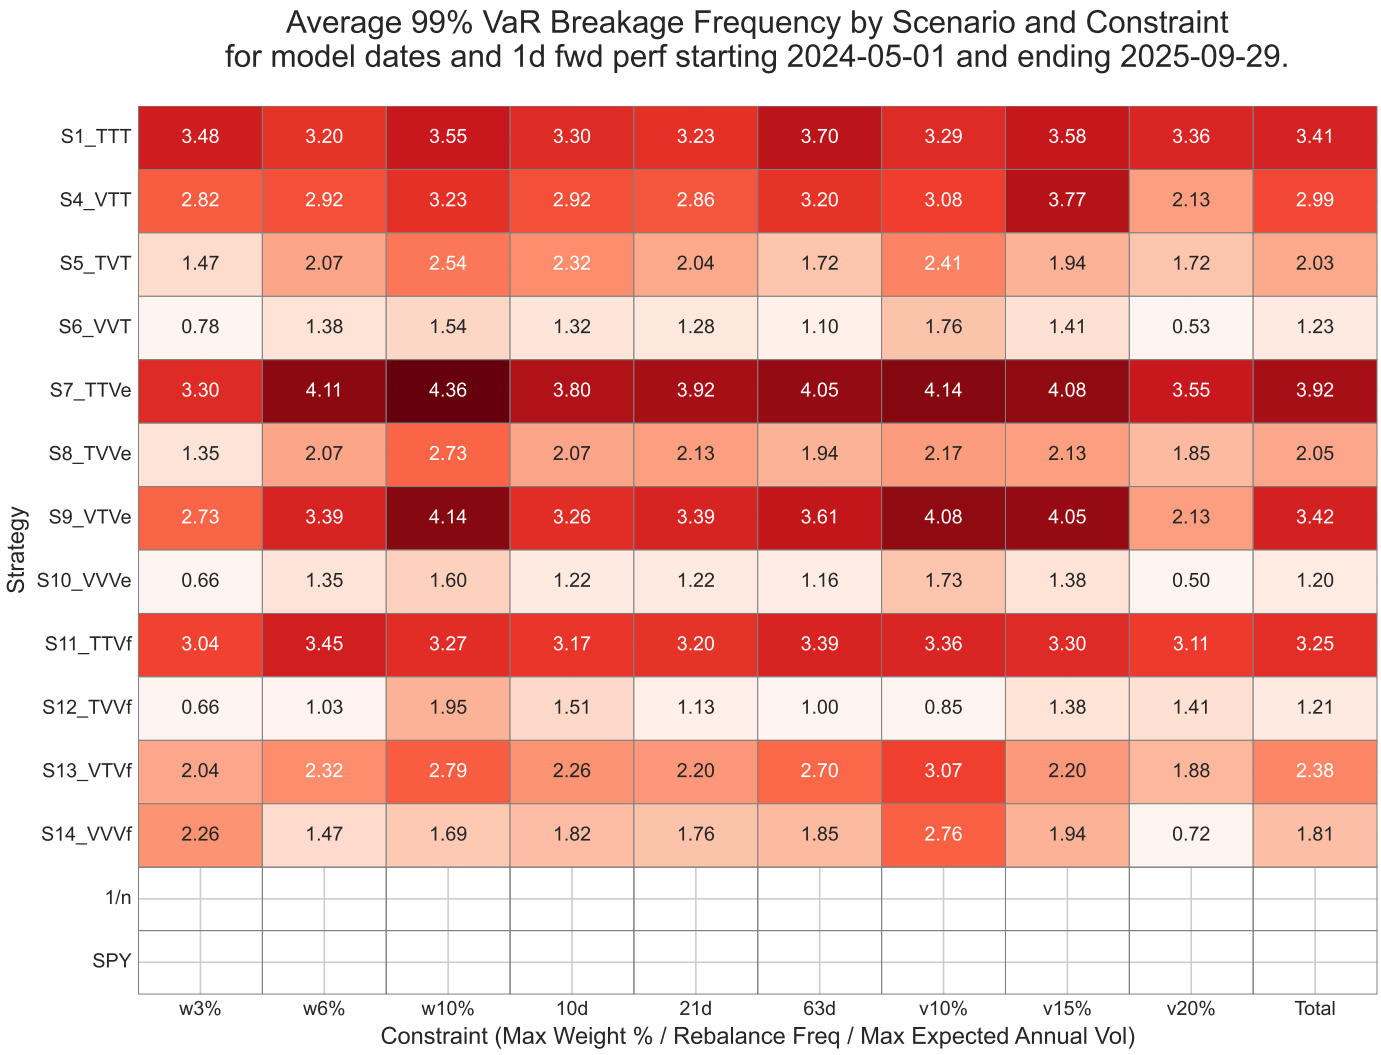

*### 99% VaR Breakage for a strategy represents the percentage frequency with which daily returns are less than -2.326*annualized Volatiltiy constraint / squrt(252). The ideal outcome would be 1.0%. Strategy returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. 99% VaR Breakage is not included in SummaryZ becasue it is well covered by Kupiec Test P-Value. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.*

Appendix: Christoferson P-Value

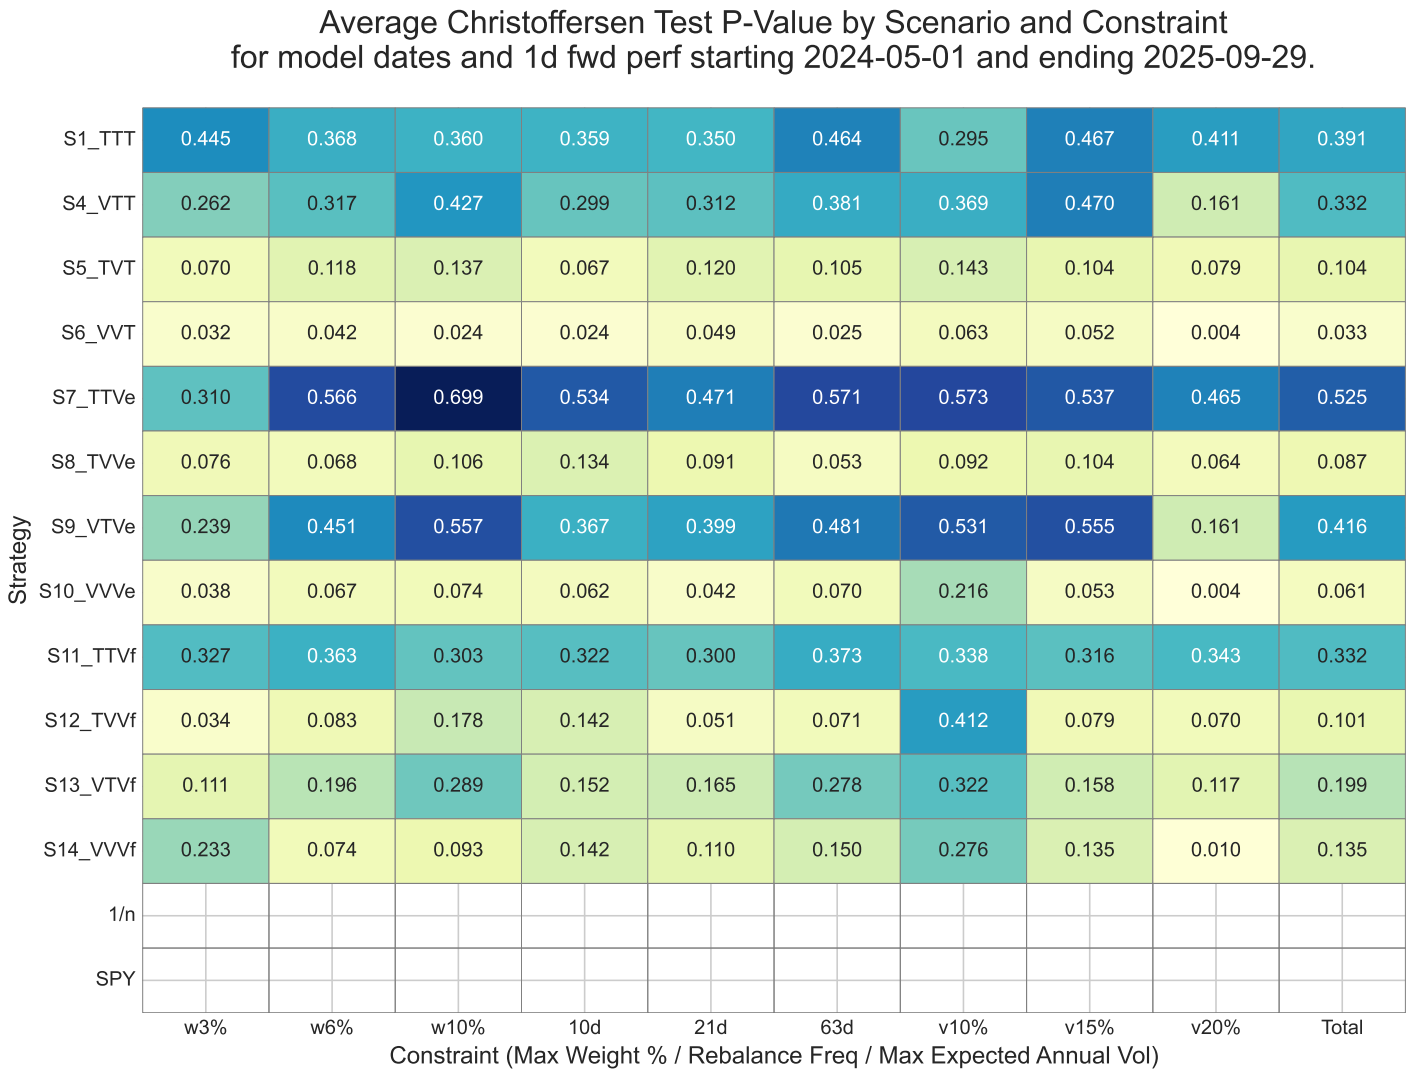

The P-Value of the Christoferson VaR Violation Independence test statistic represents the probability that the null hypothesis that the VaR model violations are independent is true (higher Christoferson P-Value is better). Christoferson is not included in SummaryZ, as its information is well represented by Max DrawDown. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Realized Vol MAE to Vol Constraint

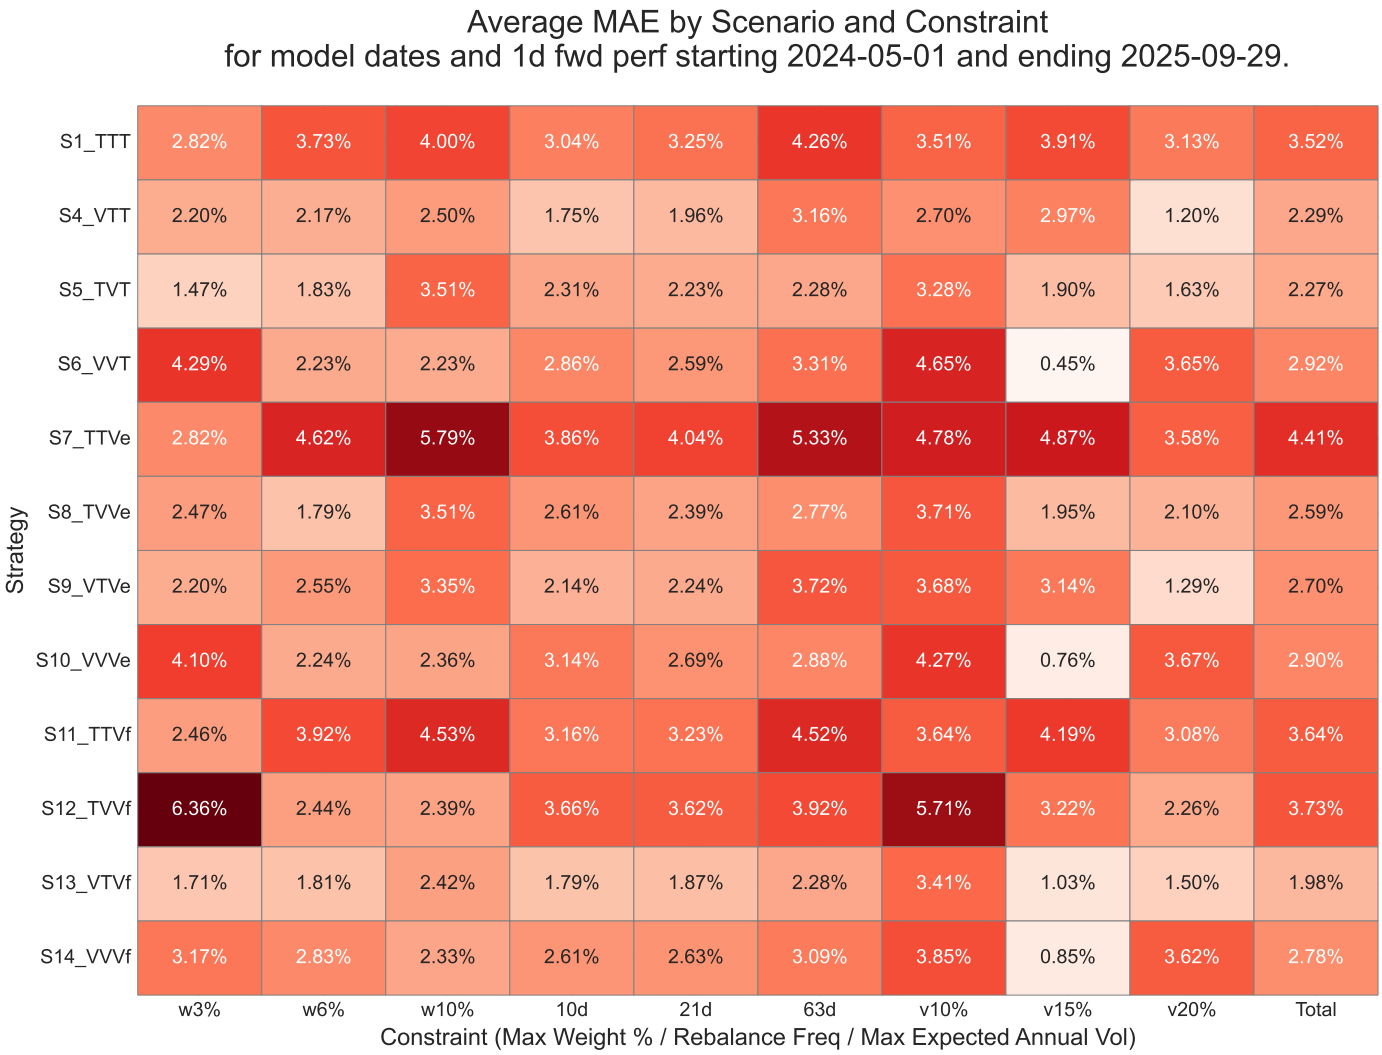

The Mean Absolute Error (MAE) relative to Vol Constraint is the MAE of realized strategy vol to the vol constraint (lower is better). Realized Vol MAE is not incuded in SummaryZ. It provides similar information as Kupiec P-Value. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Rebalance Date 2-Way Turnover

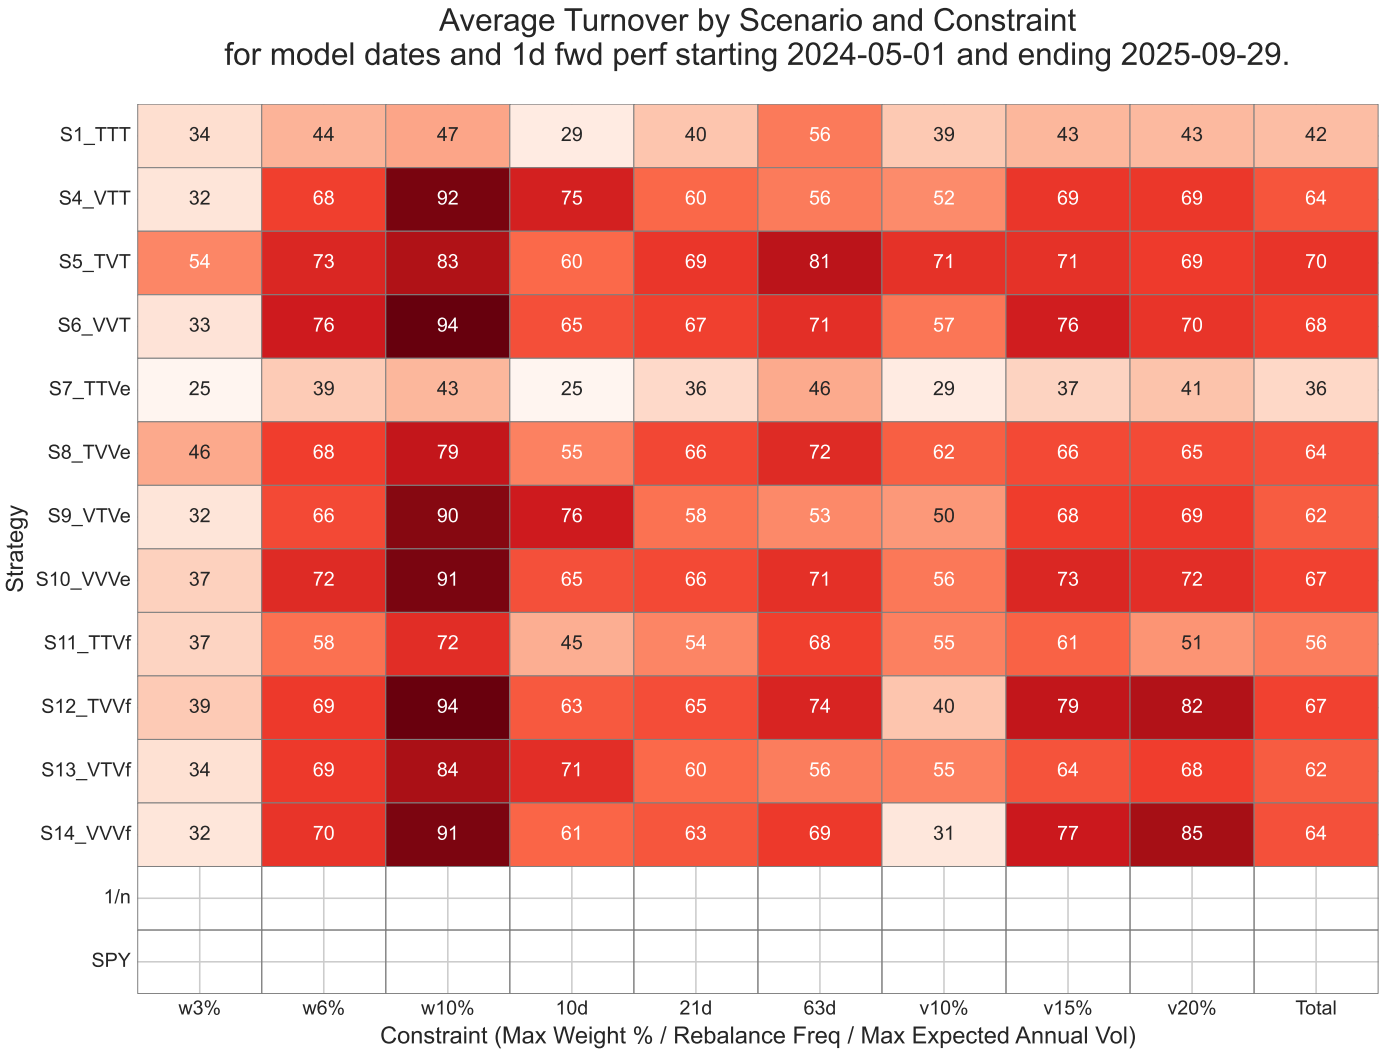

Turnover is the percentage of the portfolio the strategy buys and sells each rebalance date. For example, if 25 percent of the portfolio was bought and sold on the average reblance date then the turnover would be 50 percent. Note that these turnover statistics exclude the turnover associated with daily rebalancing to the most previously furnished set of optimized ticker weights between scheduled rebalance dates that is implicit in the return calculations for each strategy presented, including "1/n". ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Top Ticker Exposure Analysis

Top 35 Most Popular Tickers - Average Allocation by Strategy

Ticker	S1	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14
ORLY	2.92	3.16	3.48	3.36	2.30	3.14	3.04	3.28	2.33	2.72	3.94	3.84
GLD	4.35	2.52	1.39	2.47	5.33	3.32	2.22	3.10	3.62	0.96	3.13	2.64
TDG	1.25	3.56	2.72	3.63	1.16	2.51	3.62	3.63	1.56	2.55	4.43	4.34
NFLX	4.37	2.71	2.63	2.25	3.56	2.17	2.72	2.01	4.05	2.00	1.90	1.80
TMUS	4.57	2.66	2.40	2.58	3.69	1.13	2.81	2.37	2.80	1.02	2.48	2.08
QQQ	0.06	4.00	0.34	1.46	3.38	1.49	5.37	3.14	1.35	1.23	5.66	2.86
JPM	1.73	4.32	0.86	2.38	2.80	1.31	4.81	2.81	1.32	0.96	4.39	2.53
VCSH	2.09	0.66	3.70	2.02	1.92	4.29	0.65	2.53	2.66	4.23	1.31	2.99
COST	3.32	4.80	0.18	1.00	1.75	0.08	4.67	0.88	2.98	0.42	5.64	2.07
MSTR	3.44	1.01	4.37	1.34	3.69	4.71	0.96	0.80	3.30	3.42	0.34	0.36
SPY	0.00	4.17	0.15	1.07	2.08	0.70	5.12	2.63	1.49	0.53	5.62	3.46
MSI	2.48	3.80	0.25	1.17	3.12	0.65	4.32	2.02	2.58	0.51	4.00	1.93
VST	4.94	0.84	2.11	0.26	5.66	2.73	1.30	0.67	4.39	1.69	0.49	0.72
AZO	1.69	3.32	1.17	3.64	0.61	0.52	2.93	3.13	0.81	0.72	3.30	3.92
GILD	3.03	1.15	3.18	2.00	2.75	3.28	1.01	2.06	2.13	1.46	0.87	1.52
META	1.95	1.12	2.58	2.08	1.73	2.52	1.29	2.22	2.66	1.89	1.38	2.25
ORCL	0.77	2.62	2.12	3.04	0.74	1.30	3.10	3.36	0.66	0.77	1.90	2.55
T	5.03	0.14	3.64	0.30	3.03	2.10	0.10	0.10	4.60	2.58	0.16	0.34
MSFT	0.22	1.70	0.87	2.48	0.73	1.81	2.61	4.10	0.47	1.04	2.83	3.20
GWW	0.61	2.60	1.49	2.71	0.12	0.42	2.23	2.01	0.79	1.16	3.12	2.63
ISRG	0.98	2.24	0.37	0.78	3.87	1.18	3.11	1.56	1.91	0.53	2.06	1.25
MUB	0.93	0.01	2.71	1.71					1.45	3.39	0.31	2.46
LLY	2.16	2.29	1.15	1.69	2.18	0.74	2.44	1.35	1.44	0.62	1.51	1.64
POST	1.65	1.49	2.95	2.91	0.69	2.18	1.30	2.97	0.07	0.36	0.82	1.62
GOOGL	0.85	3.09	1.28	2.11	0.24	0.78	2.41	1.18	0.57	1.02	2.73	2.76
WFC	2.01	0.52	2.86	1.96	1.66	3.43	0.45	2.15	0.46	0.72	0.42	1.88
GS	0.67	1.86	0.33	0.87	3.39	1.21	2.56	1.49	1.73	0.50	2.07	1.86
ACGL	1.18	3.47	1.46	2.27	0.08	0.15	3.07	1.27	0.11	0.29	3.05	2.04
HCA	1.16	1.59	2.14	2.31	0.83	1.71	1.68	2.67	0.48	0.69	1.11	1.90
SLV	1.01	0.26	2.64	1.03	3.76	4.60	0.39	1.55	0.58	1.28	0.10	1.02
AVGO	1.80	2.36	0.94	0.51	2.93	0.95	2.70	0.50	2.07	0.92	1.53	0.44
PWR	0.57	3.18	0.85	2.40	0.37	0.42	2.83	1.23	0.58	0.82	2.36	1.78
TEVA	3.50	0.21	1.91	0.18	3.66	2.16	0.25	0.25	3.12	1.38	0.10	0.23
CDNS	0.11	1.29	1.57	2.95	0.09	0.96	1.39	2.65	0.11	1.34	1.28	2.87
NVDA	2.83	1.87	0.27	0.23	2.46	0.25	1.94	0.20	3.00	0.43	1.15	0.30

Strategy

Appendix: Bottom Ticker Exposure Analysis

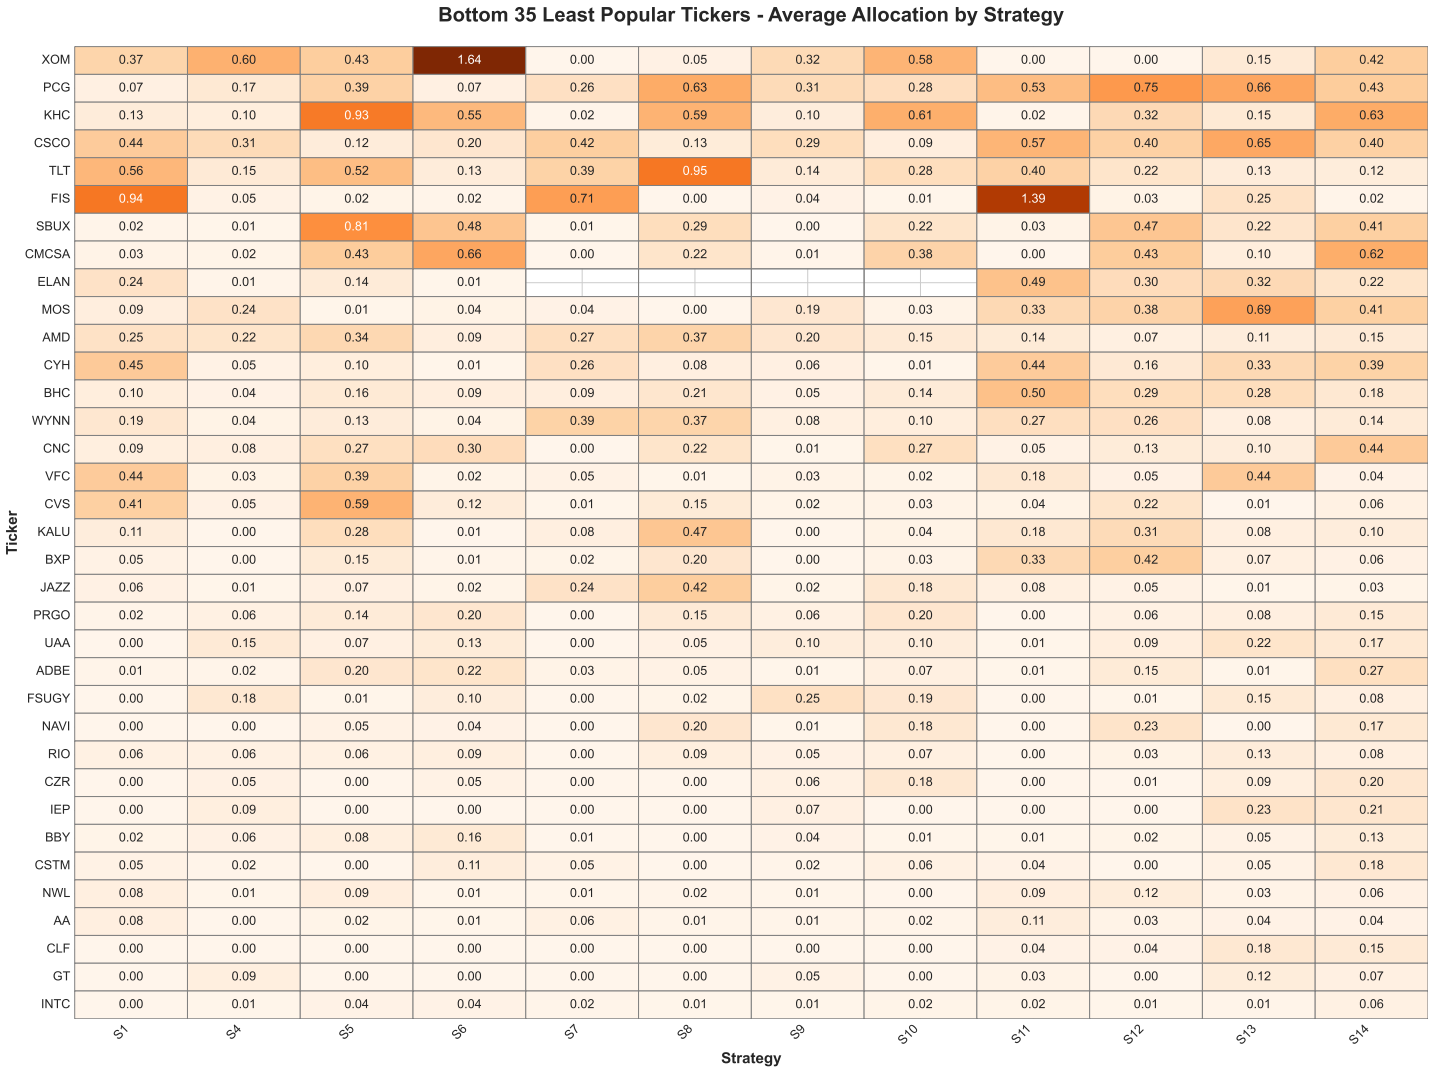

Appendix: Comparison of Correlation Distributions Across Pearson/ VecEvent/ Fingerprint

Here are the distributions of pairwise correlation values generated by the three correlation methodologies that are utilized in the grid search during the Training Period, which spans 1/31/2022 through 4/30/2024. For Pearson we use a trailing 252d window, just as we do in the grid search. Thus, the distribution for Pearson represents trailing 252d correlations values from 1/31/2023 through 4/30/2024. "Fingerprint" correlation covers the entire period on a daily basis, spanning 1/31/22 through 4/30/2024. Finally, "VecEvent" based correlation is kept static throughout, given the point in time ambiguity of LLM responses concerning the past.

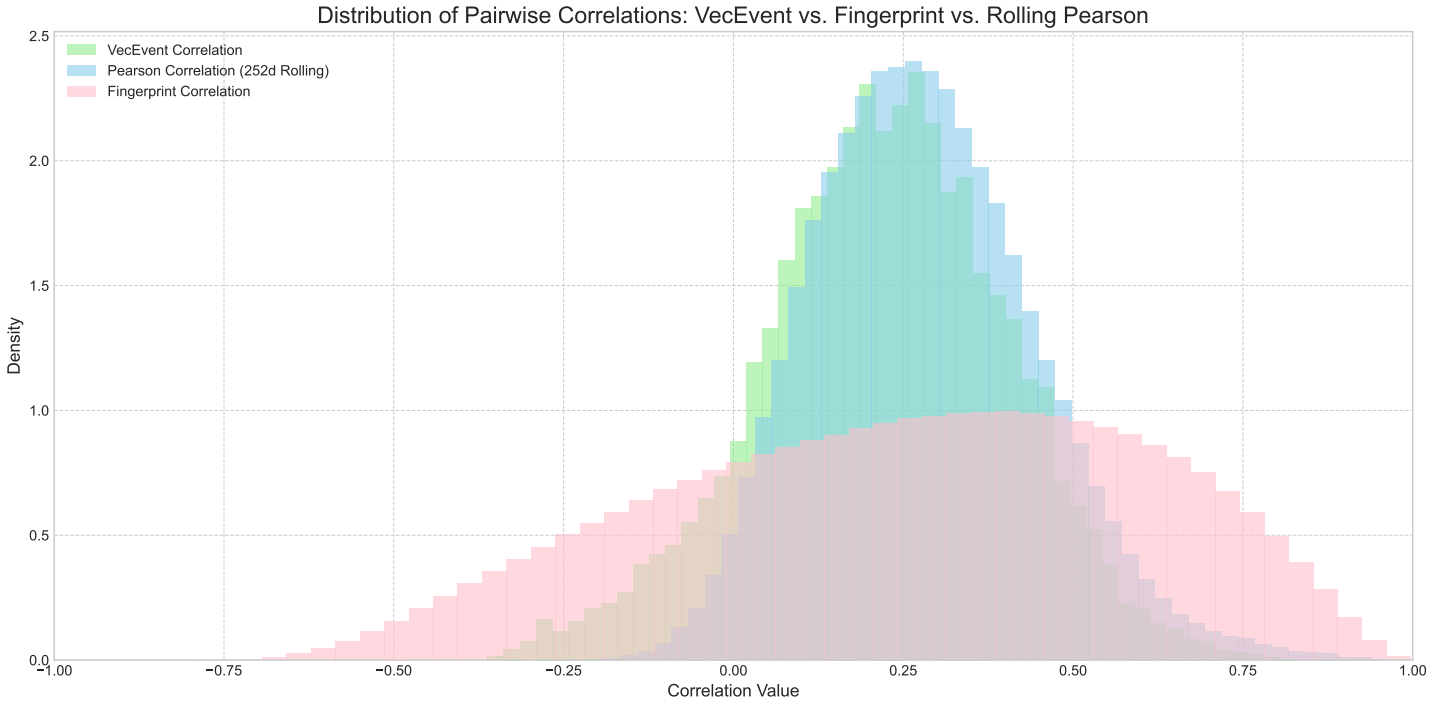

Appendix: VecEvent Similarity Correlation Methodology

Introduction

VecEvents are brief descriptions of an event or theme thought to be influential upon a ticker. They are sourced through Gemini 2.5 Pro, a large-language model (LLM), by VecViz. VecViz links them to the Vector Set channels that are the core building block of its machine learning based Vector Model.

Linkage of a VecEvent to a Vector Set is done on the basis of the overlap of the dates the LLM indicates the VecEvent was influential upon a ticker with the dates of the major price tops and bottoms anchoring the Vector Set. See below for an example of a VecEvent presented in conjunction with its Vector Set, taken from the Dashboard page of vecviz.com.

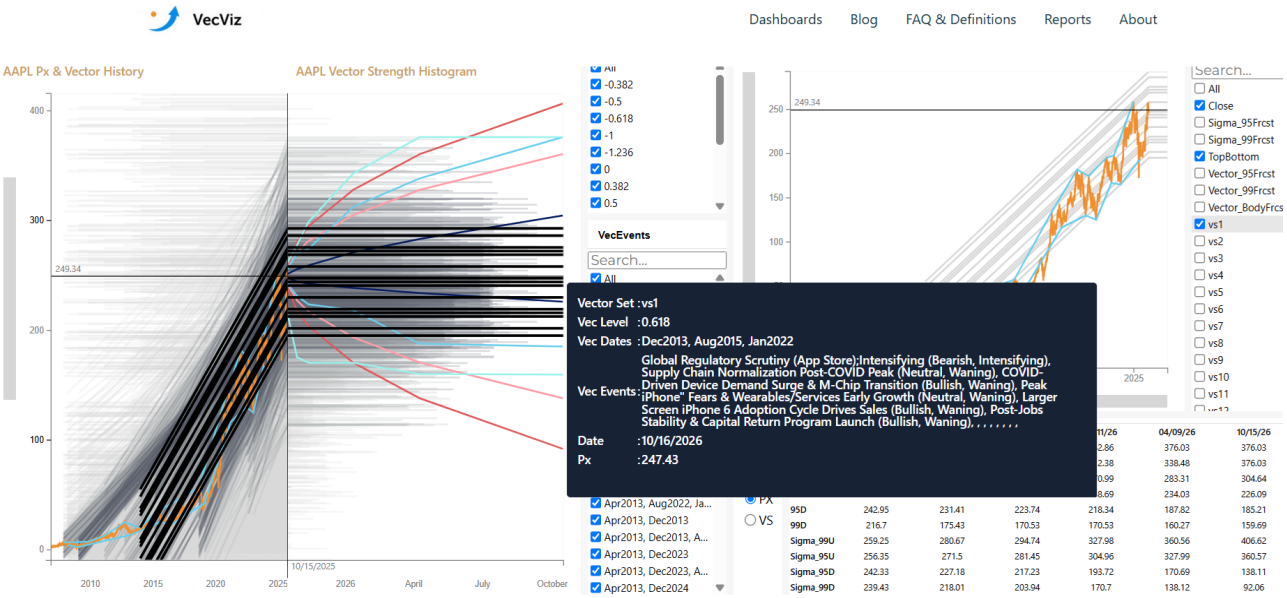

We theorize that these VecEvent narratives capture the key drivers of ticker performance. As Ding et al. (2014) note, "Events reported in financial news are important evidence to stock price movement prediction", suggesting that structured event representations can inform return forecasts.

Building on this insight, we hypothesize that VecEvent similarity will correspond to price behavior similarity (i.e., correlation of price returns). Here we detail how we created numerical features from VecEvent text that reflect price co-movement between tickers and are applicable to MVO and perhaps other portfolio optimization methodologies as well.

Feature Engineering

To analyze the unstructured text of corporate events, we employ a multi-stage methodology beginning with numerical representation. The initial step uses sklearn's TF-IDF (Term Frequency-Inverse Document Frequency) vectorizer function to convert each event description into a meaningful numerical vector.

TF-IDF starts by building a comprehensive vocabulary of unique terms from the entire corpus after filtering for common stop words. Each event's text is then transformed into a vector where each position corresponds to a word in the vocabulary, and the value at that position is a weighted score. This score is highest for terms that appear frequently within a single document but are rare across all other documents, effectively highlighting significant and descriptive words like "acquisition" or "FDA-approval."

Once every event is represented as a high-dimensional vector, we distill these thousands of data points into thematic groups using sklearn's K-Means clustering function. K-Means identifies distinct clusters of similar events and calculates their central vectors, or centroids. These centroids serve as our analytical "benchmarks," with each one representing the numerical profile of a common event category.

We then aggregate from the event level to the firm level by creating a single summary profile for each unique company. This is achieved by grouping all event vectors for a given ticker and creating a composite vector by taking the position-wise maximum value. This method creates a profile that captures the peak significance of the most defining terms ever used in a company's events, rather than an average which could dilute the impact of rare but critical occurrences.

With a unique profile vector for each company, we then measure the alignment of each company to our established themes. Using sklearn's cosine similarity function, we calculate a score from 0 to 1 that quantifies how closely each company's summary profile matches each of the benchmarks.

The result is a benchmark exposure profile for every firm, which serves as a unique signature of its activities. From these profiles, we can calculate the correlation between any pair of companies, revealing which firms exhibit similar patterns of behavior based on their event histories. We utilized Pearson correlation instead of cosine similarity at this point because it gives considers variation from the mean. Note that by anchoring exposure to each ticker's mean benchmark exposure we are attempting to control for bias in the benchmark formulations toward industries that are more highly represented in the tickers included in this study (e.g., tech, metals, pharma, banking). These correlations can be useful in MVO style portfolio optimization.

Finally, to ensure the interpretability of our abstract benchmarks, we assign a human-readable label to each one. This is done by identifying the single, original event description from the dataset that has the highest cosine similarity to each benchmark vector, allowing this real-world example to define the theme of its corresponding cluster.

This approach is computationally efficient, produces interpretable benchmark archetypes that analysts can review, and generates structured features directly usable in portfolio construction methods such as MVO.

VecEvent Benchmark Optimization Process and Results

To identify the optimal number of VecEvent benchmark clusters, we used forward returns from 1/31/2022 to 4/30/2024 together with all VecEvents to determine the correlation matrix that minimized error against market returns. We then "tested" this configuration by checking its performance against forward returns from 4/30/2024 to 7/9/2025. Benchmark counts were varied from 2 to 201, with Mean Absolute Error (MAE) used as the comparison metric.

Key Findings:

1) Optimal Benchmarks: Using the full ticker set, the optimal configuration was 192 clusters, yielding a minimum MAE of 0.155165.
2) Out-of-Sample Performance: When tested against the forward return correlations from 4/30/2024 to 7/9/2025, the same 192-benchmark correlation produced an MAE of 0.155173, only slightly higher than the training MAE.
3) Subset Analysis: To evaluate the sensitivity of results to the number of tickers, we ran randomized subsets of sizes ranging from 30 to 138, evaluating MAE vs. actual Pearson correlation in the Training data sets.
4) These runs suggest that as the ticker universe grows, the optimal benchmark count tends to rise, though the differences in MAE remain relatively modest.
5) Ordering Sensitivity: In some runs, the optimal benchmark number fluctuated slightly lower as ticker count increased depending on dataset ordering. This effect may be due to KMeans sensitivity to input order or hidden inconsistencies in event text formatting. Sorting the dataset by Ticker and Start Date reduced this variability and produced stable results.

We display the simulation results from our testing of 138 tickers below:

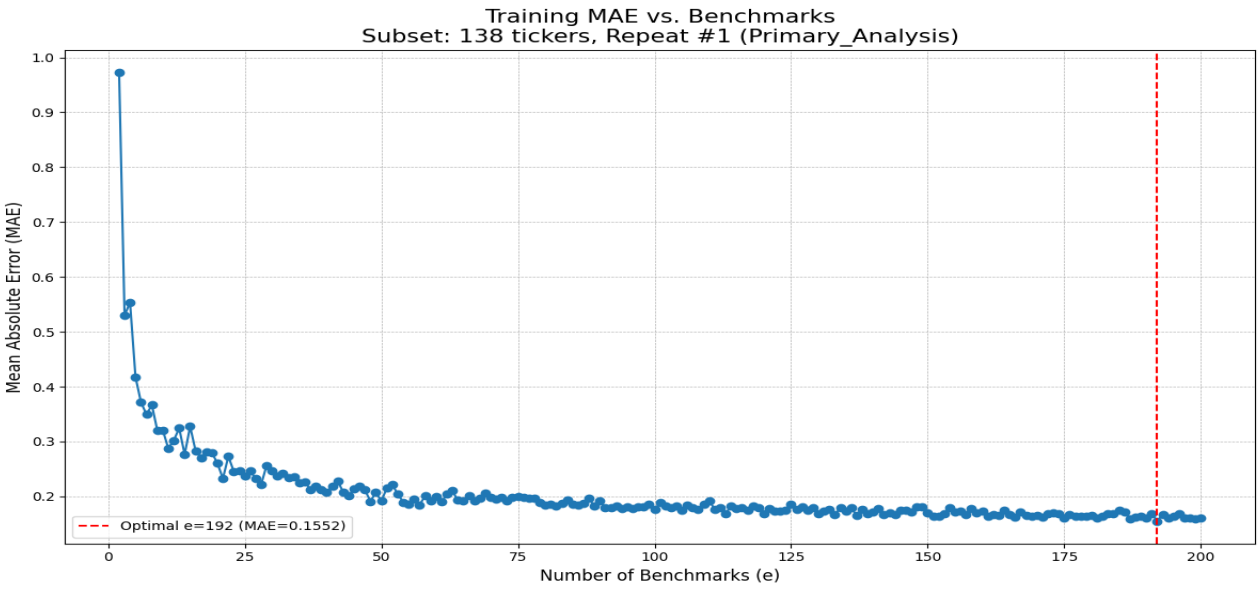

Interpretation

The VecEvent-based correlations capture real-world co-movement structures effectively at the ~192-benchmark level, with strong out-of-sample consistency. Subset analysis shows robustness across different ticker universes, though optimal benchmark counts trend upward with larger sets. For HRP applications, a reduced number of benchmarks at an "elbow point" remains a practical option, balancing model manageability with performance.

Alternative Approaches, Bias & Limitations

1) While this study relies on TF-IDF vectorization and KMeans clustering of VecEvents, several extensions could improve robustness and interpretability:

   a) Regular VecEvent Updates: Future work should implement a regular cycle of refreshing VecEvents, dating each batch by retrieval day, and re-measuring both implied correlations and their error against actual forward return correlations across multiple horizons. This would help detect and mitigate any look-ahead bias introduced by the LLM sourcing process.

   b) Enhanced Feature Representations: Transformer-based embeddings or other semantic encoders could capture richer contextual meaning beyond bag-of-words statistics, though at the cost of higher computation and reduced transparency.

   c) Alternative Clustering Methods: Hierarchical or spectral clustering might produce benchmark groupings better aligned with HRP's tree-based allocation logic, or reveal multi-scale structures hidden by KMeans.

2) Observed Data Sensitivities:

   a) Ordering Effects: Results showed slight variation in the optimal benchmark count (e.g., 183 vs. 187) depending on dataset ordering. Sorting events by Ticker and Start Date produced stable outputs, suggesting KMeans sensitivity to random initialization and hidden text artifacts.
   b) Subset Size Dependence: Optimal benchmark counts tended to rise as the ticker universe grew, though MAE differences remained small. This implies that correlation structures scale with market breadth but remain stable in predictive accuracy.

3) Bias Sources:

   a) LLM Event Generation Bias: VecEvents are created by a large-language model, inheriting any training-data or prompt biases. This could lead to selective emphasis or omission of certain event types or industries.
   b) Historical Coverage Bias: The underlying financial news coverage may underrepresent smaller firms or certain sectors, influencing both event diversity and correlation estimates.
   c) ETF vs. Single-Ticker Asymmetry: By design, ETF VecEvents are macro-oriented, whereas single-company events are micro-specific. This leads to systematically lower VecEvent-based correlations between ETFs and individual tickers. A possible remedy for equity ETFs is to compute a market-cap-weighted composite of their constituent ticker correlations; for bond ETFs, adjustments may be more limited.
   d) When generating VecEvents in a long term retrospective manner it is never clear when the LLM (or a human fundamental analyst) would have acknowledged the VecEvent's inception, limiting the confidence one can put in backtests relying upon such VecEvents.

4) Limitations:

   a) Static Clustering: Once KMeans clusters are trained, they remain fixed and do not adapt to evolving event patterns or market regimes.
   b) Linear Correlation Assumption: The approach assumes that textual similarity translates linearly into return correlations, which may not hold in all market conditions.
   c) Data Volume and Horizon: Filtering VecEvents to pre-2024 influence start dates reduces look-ahead bias but doesn't conclusively eliminate it.

Appendix: Fingerprint Similarity Based Correlation Methodology:

"Fingerprint" Similarity considers similarity in standardized Vector Model and V-Score input and output metrics.

Instead of using historical price returns to calculate correlation, this code computes correlation based on the similarity of 20 features that together comprise a ticker's "fingerprint", from a VecViz perspective. It does not utilize VecEvent information (at this point).

On each Model Date, each of the 20 criteria for each ticker is converted into a percentile ranking. Fingerprint similarity for a ticker pair is simply the Pearson correlation of their values for those 20 standardized criteria. The thesis is that if two stocks have very similar feature values on a given day, (e.g., both are high momentum, low value), they will have a high positive forward correlation, and vice versa. This expectation is supported by the use of forward looking, machine learning generated analytics in the fingerprint criteria. There was no training or calibration of this methodology.

All these pairwise "similarity scores" are then assembled into a standard correlation matrix, which can then be used in a portfolio optimization process.

Here are the 20 features that together constitute the "fingerprint": 1) There are 6 Vector Chart Shape Features. These metrics are inputs to VecViz's Vector Model, and V-Score:

PR_RangeUp PR_RangeDwn PR_S_VStr/R_VStr PR_Wgt Up PR_DaysSinceLastTop PR_LevelUpPctR

The "PR_" preceding each variable name refers to "percentile rank". Explanations for the variables listed here can be obtained on vecviz.com, on the dashboard page, by hovering your mouse over their name on the V-Score spider chart. See pic below.

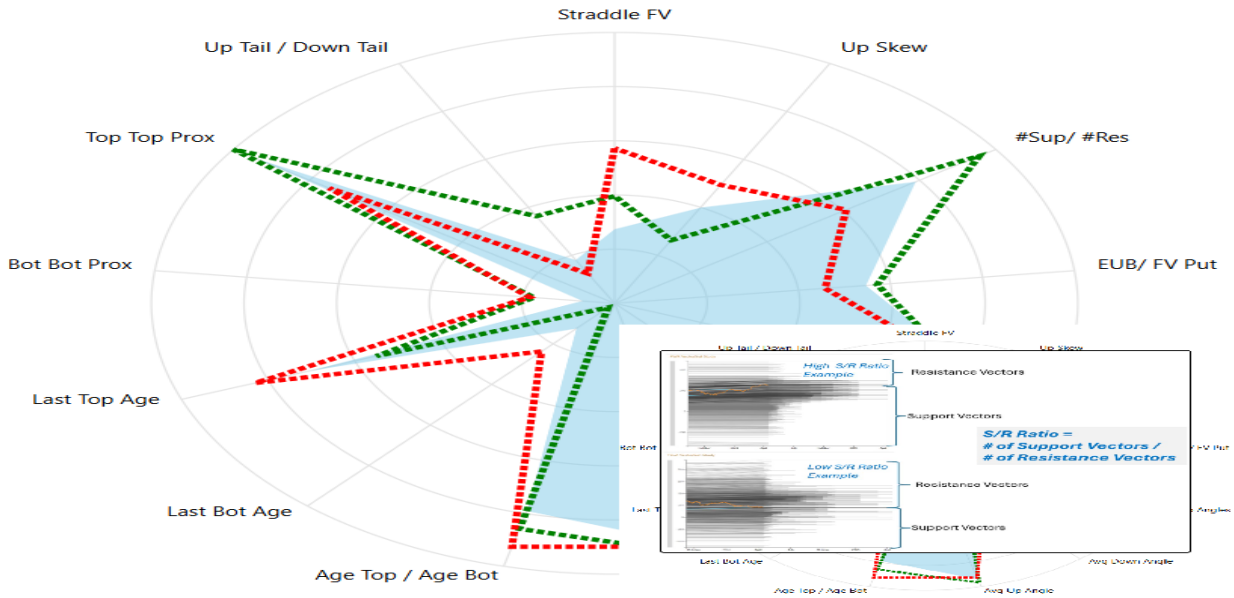

2. There are 7 Vector Model Probability Features:

PR_95D_Ret PR_EUB_Sratio PR_EDB_Sratio PR_95U_Sratio PR_95D_Sratio PR_99U_Sratio PR_99D_Sratio

With the exception of "PR_95D_Ret", each of these look at the ratio of a Vector Model price probability percentile relative to a corresponding price probability percentile from a variation of "Trailing". This variation of "Trailing" is a common approach to price volatility, which weights daily return observations in the volatility calculation lookback window with exponential time decay. This metric is presented alongside Vector Model output in many reports and dashboards on vecviz.com, where we refer to it as "Sigma".

Some explanation of the abbreviations: "95" and "99" refer to the 95th and 99th percentile, respectively. "U" and "D" refer to upward and downward, respectively. "EUB" and "EDB" refer to "Expected Up Body" and "Expected Down Body", the probability weighted averages between the Model Date price and the 95U and 95D prices, respectively.

The "SRatio" is calculated to represent the ratio of the position size an investor would take if using the Vector Model to forecast price volatility relative to the position they would take if using Sigma, subject to a cap and floor at 3.0 and 0.333, respectively. Thus, if the Vector Model says 95D for a stock is -30% and Sigma says 95D is -20%, then the 95D_SRatio for the risk adverse, VaR aware investor considering that ticker would be 0.67 (i.e., -20/-30). Likewise if the 99% Vector Model OaR for a ticker is +24% and the 99% Sigma OaR for that ticker is +16% then the 99U_SRatio for the return seeking, OaR aware investor considering that ticker would be 1.5 (i.e., 24/16).

3. There are eight V-Score Expected Return Composite Features: PR_V-Score (PR_VS+RU+SR+WU-RD-LU-DSLT)/95D (PR_VS+RU+SR+WU-RD-LU-DSLT)/7 PR_VS_NetAdj PR_VS_AddAdj PR_VS_DownAdj PR_VS_UpAdj PR_Blend

These features are combinations of the features listed in the first two groupings, along with the percentile rank ("PR_") for the V-Score. They are described in the VecViz Expected Return Composite Section that follows.

Appendix: VecViz Regime Based Expected Return Composites

VecViz's V-Score, is included in all composites. As discussed earlier, the V-Score's training data ends on 1/31/2022, capturing a full 1 year forward horizon of the last V-Score metric trained, which was calculated on 1/31/2021.

For purposes of this study, and in recognition of the fact that the V-Score could benefit from a refresh, we supplement the V-Score with regime based "top-side" adjustments. These adjustments are done by standardizing the V-Score on a percentile scale and adding or subtracting similarly standardized percentiles of related inputs to or from it.

There are eight V-Score Expected Return Composite Features. They are defined below.

1) PR_V-Score = the simple percentile ranking of the V-Score itself

2) (PR_VS+RU+SR+WU-RD-LU-DSLT)/95D = the sum of the percentiles for the V-Score (VS) and upward trajectory oriented metrics minus the percentiles for downward sloping trajectory metrics divided by the absolute value of the 95D_Ret.

The upward sloping metrics include: RU = RangeUp = the current price as a % of the highest identified top, SR= support vector strength / resistance vector strength, WU= % of total vector strength in upward sloping vector sets

The downward sloping metrics include: RD = RangeDown = the current price as a % of the lowest identified bottom, LU= the percent of resistance vector count comprised of "leveled up" vectors DSLT = days since last top

3) (PR_VS+RU+SR+WU-RD-LU-DSLT)/7 = same as item #2 but divided by 7 instead of by 95D

4) PR_VS_NetAdj = the weighted sum of the percentiles for the V-Score (VS) and "return seeking SRatios" (EUB_SRatio, 95U_SRatio, 99U_SRatio) minus the percentiles for the "risk avoidant SRatios" (EDB_SRatio, 95D_SRatio, 99D_SRatio). See "Appendix: Fingerprint Similarity Based Correlation Methodology" for discussion of SRatios.

5) PR_VS_AddAdj = the weighted sum of the percentiles for the V-Score (VS), return seeking SRatios and risk avoidant SRatios. See "Appendix: Fingerprint Similarity Based Correlation Methodology" for discussion of SRatios.

6) PR_VS_DownAdj = the weighted sum of the percentiles for the V-Score (VS) and the risk avoidant SRatios.

7) PR_VS_UpAdj = the weighted sum of the percentiles for the V-Score (VS) and the return seeking SRatios.

8) PR_Blend = average of (PR_VS+RU+SR+WU-RD-LU-DSLT)/95D and PR_VS_UpAdj

These composites and the weights utilized in PR_VS_NetAdj, PR_VS_AddAdj, PR_VS_DownAdj, PR_VS_UpAdj were identified and calibrated via exploration of the 2/28/2022 thru 4/30/2024 period. Likewise, the regimes to which they are linked were identified and calibrated via exploration of that period.

Utilization of these composites across and by regime during the 4/30/2024 through 9/30/2025 period of the study is detailed in the tables that follow below.

Appendix: Feature Utilization Across Regimes

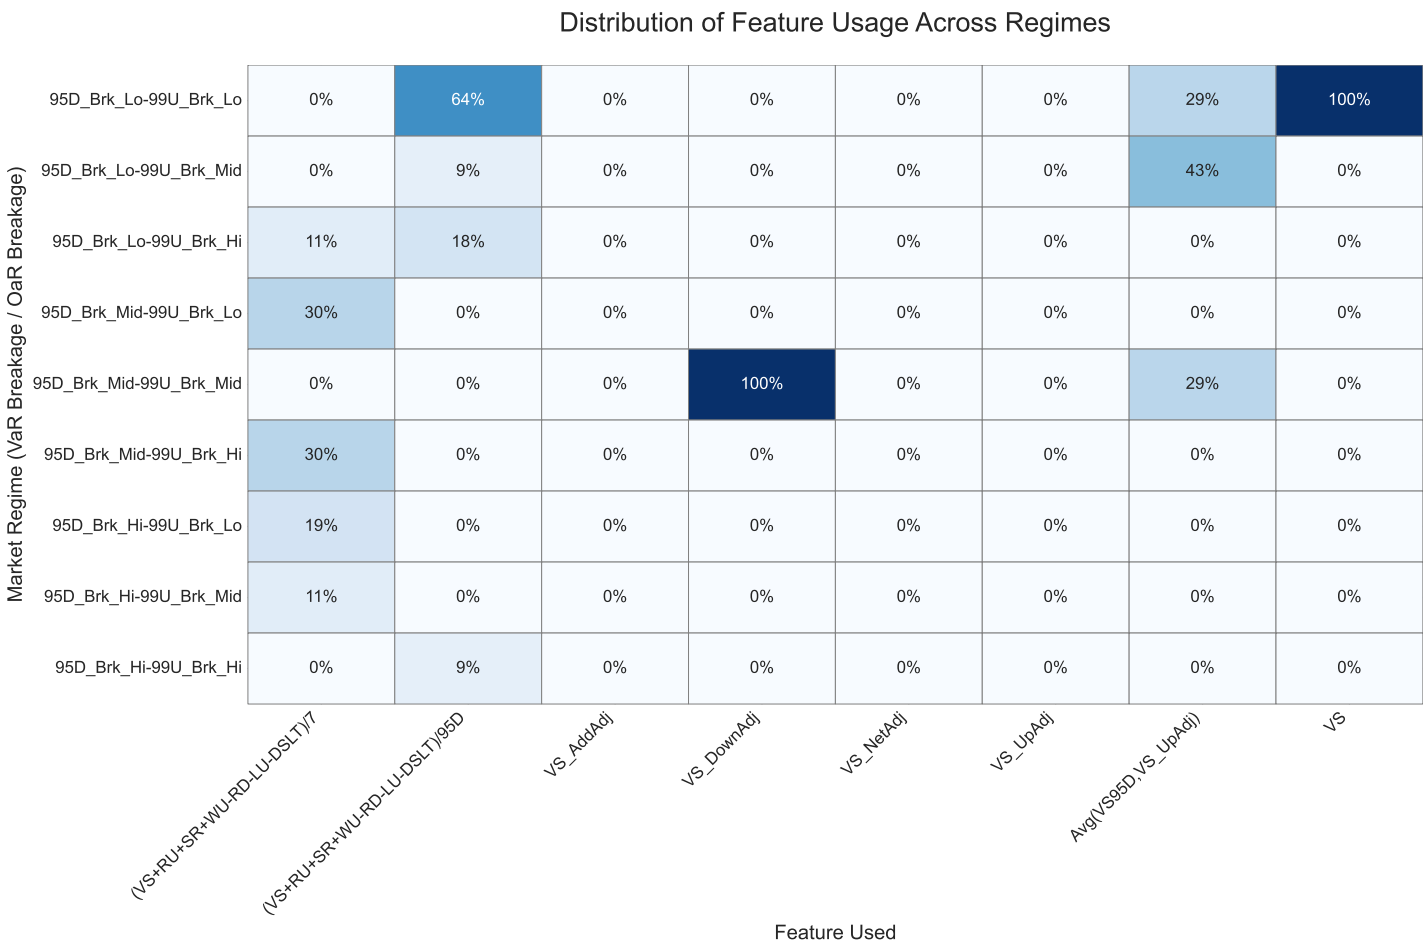

Appendix: Feature Utilization By Regime

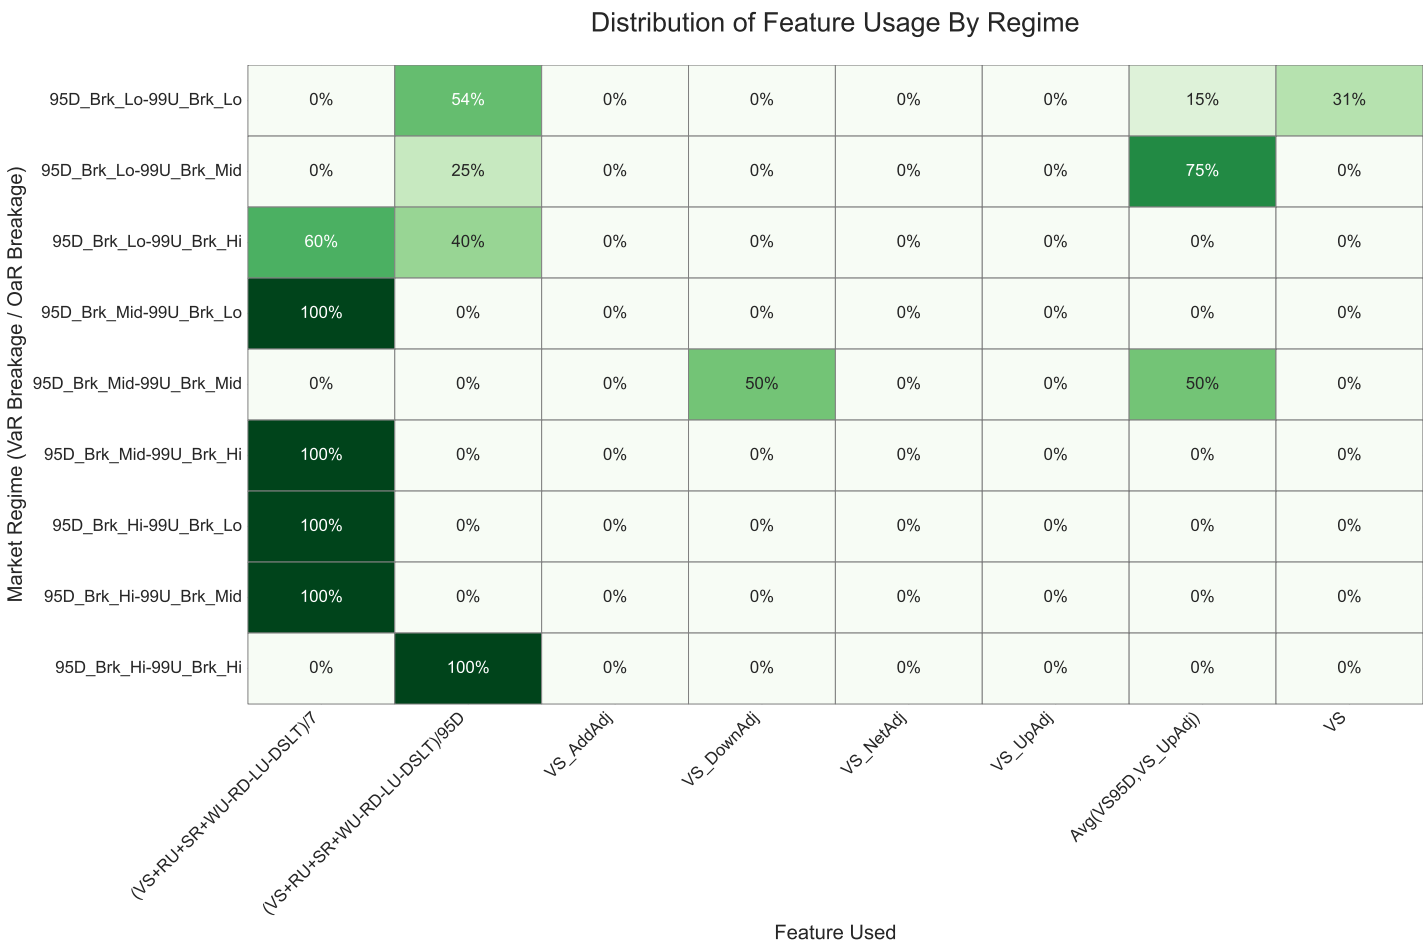

Appendix: Results For 3/31/2023 thru 9/30/2025 (includes in-sample 3/31/2023 thru 4/30/2024 period)

Appendix: Range of Ticker Performance

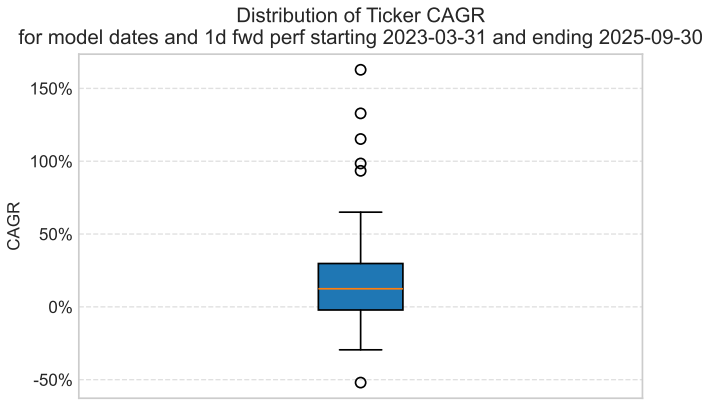

(a)

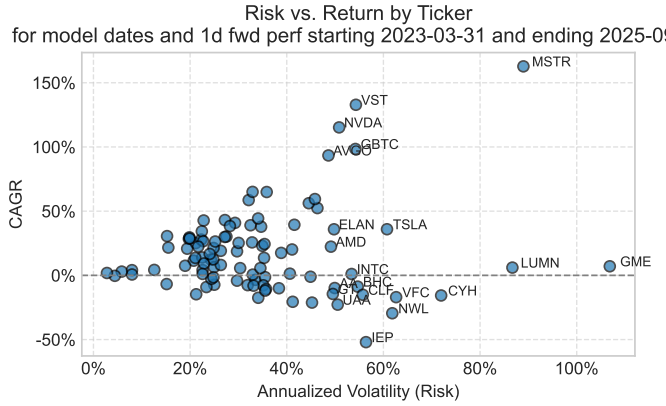

(b)

Appendix: Performance Metrics by MVO Input Variable

Average Performance Metrics by MVO Input Variable
for model dates and 1d fwd perf starting 2023-03-31 and ending 2025-09-29.

Input Variable	SummaryZ	Ann_Ret	MaxDD	Sharpe	Calmar	alpha_ann%	Kup.-Pval
Correl (VE)_VV	1.61	25.9%	-26.7%	1.65	1.02	16.5	0.24
Vol_VV	0.89	20.8%	-19.8%	1.49	1.06	8.1	0.33
Ret_T252d	0.75	26.5%	-29.5%	1.62	0.94	13.4	0.22
Correl (FP)_VV	-0.71	21.5%	-24.1%	1.46	0.92	4.0	0.24
Ret_VV	-0.76	19.5%	-21.2%	1.40	0.96	2.8	0.27
Vol_T252d	-0.89	25.2%	-30.9%	1.54	0.84	8.1	0.16
Correl_T252d	-0.91	21.6%	-25.2%	1.43	0.92	3.7	0.25

Appendix: "SummaryZ" Performance Score Table

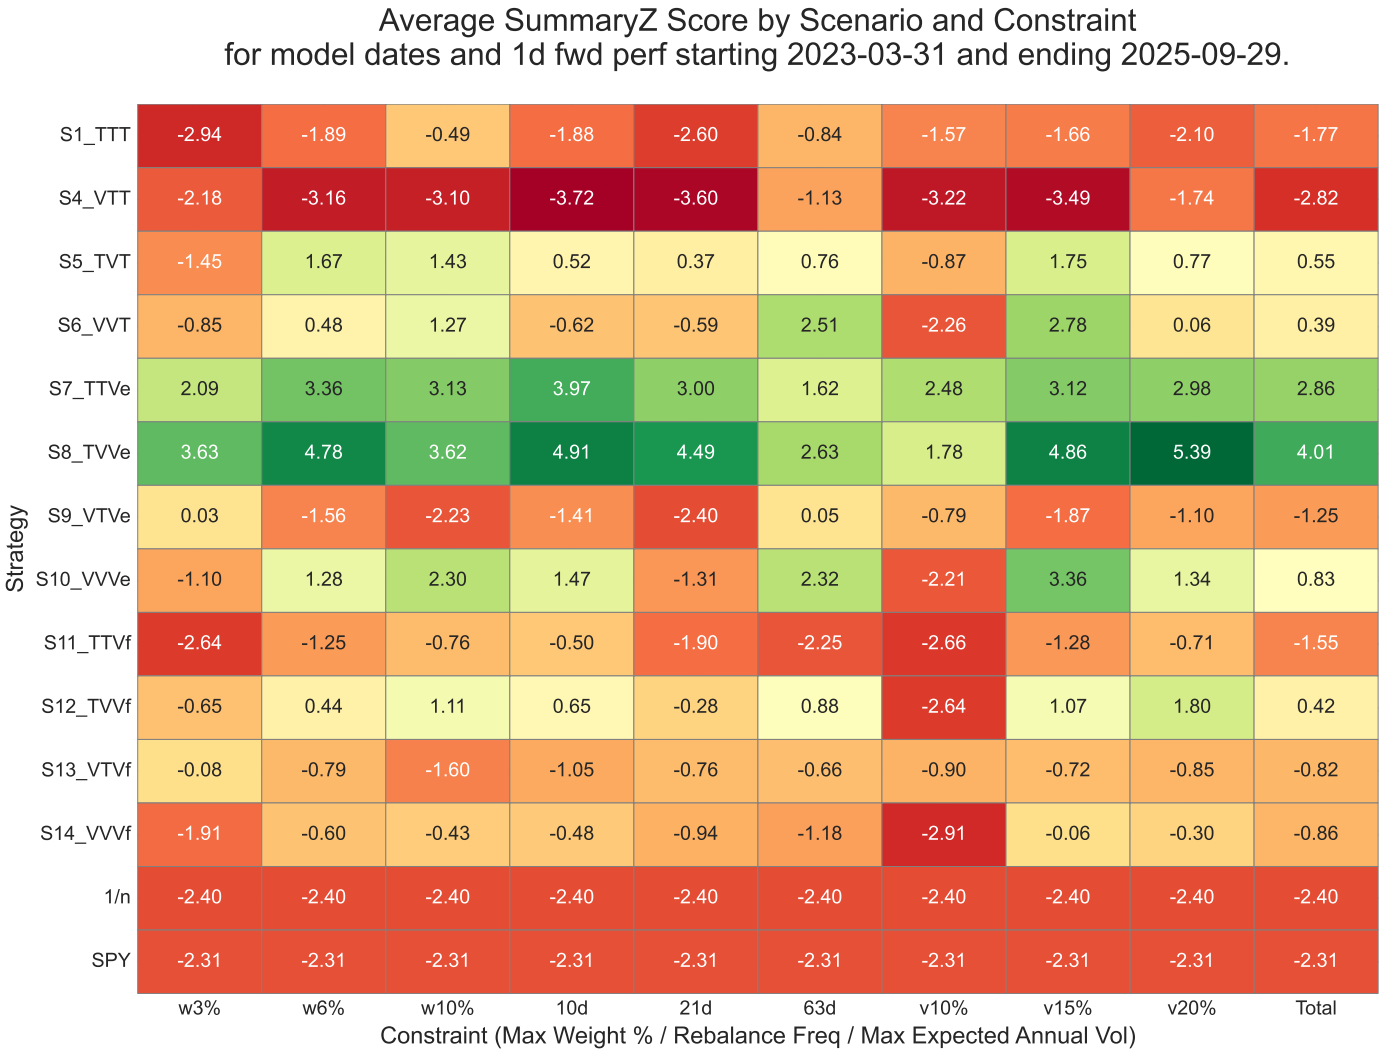

SUMMARYZ score, a sum of Z-scores calculated across all strategies and constraint combinations for six metrics: Annualized Return, Sharpe Ratio, Max Drawdown, Calmar Ratio, 99% VaR Kupiec P-Value, and Multi-Factor Alpha. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Strategy - Constraint Scenario Cumulative Return Paths

```
Generating plot...
```

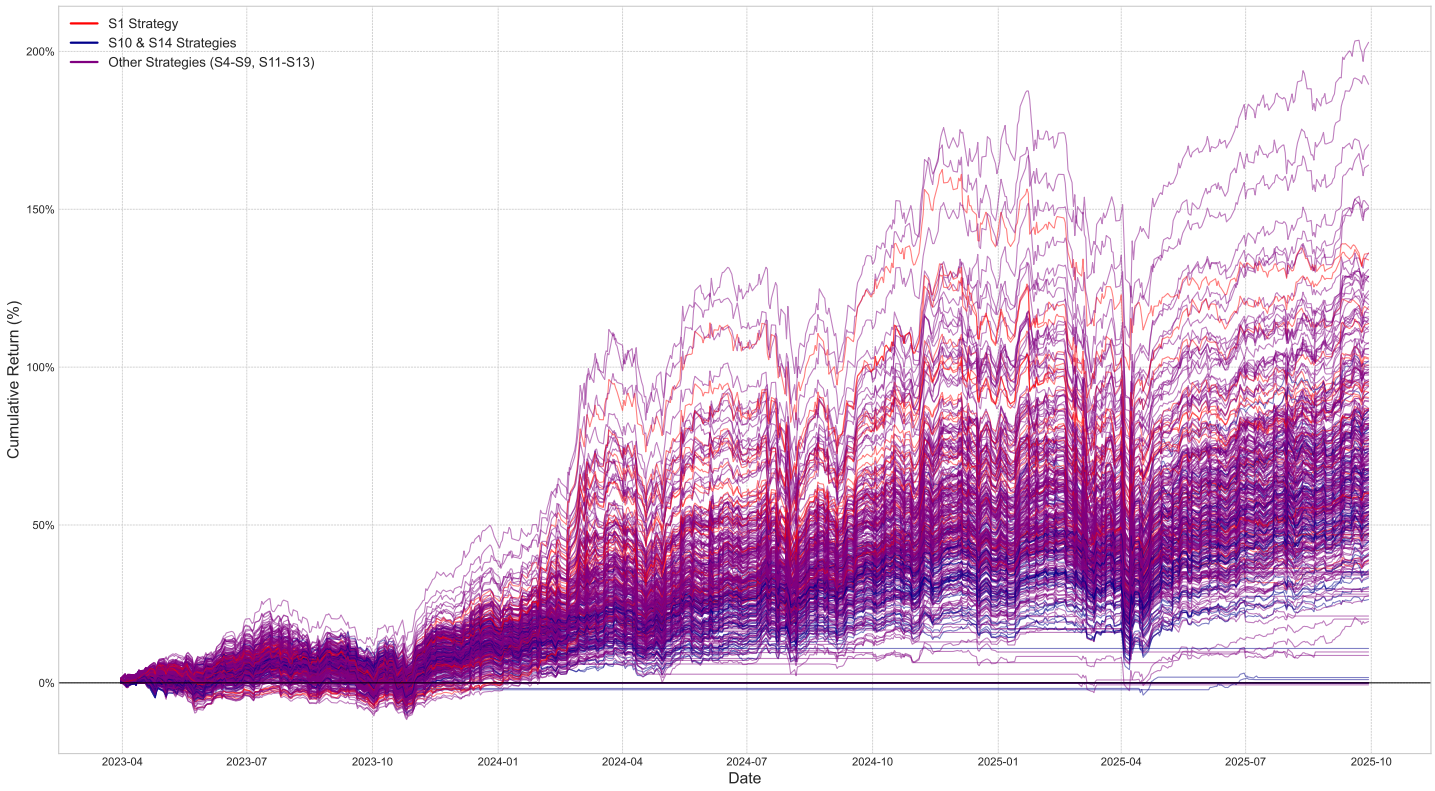

Figure 4: Cumulative Returns of All Grid Search Strategies

Appendix: Return vs. Risk Scatter Charts for 3/31/2023 thru 9/30/2025

Squares= VecViz Return, Circles=Trailing 252d Return | Red=All Trailing, Purple=Hybrid, Blue=All VecViz

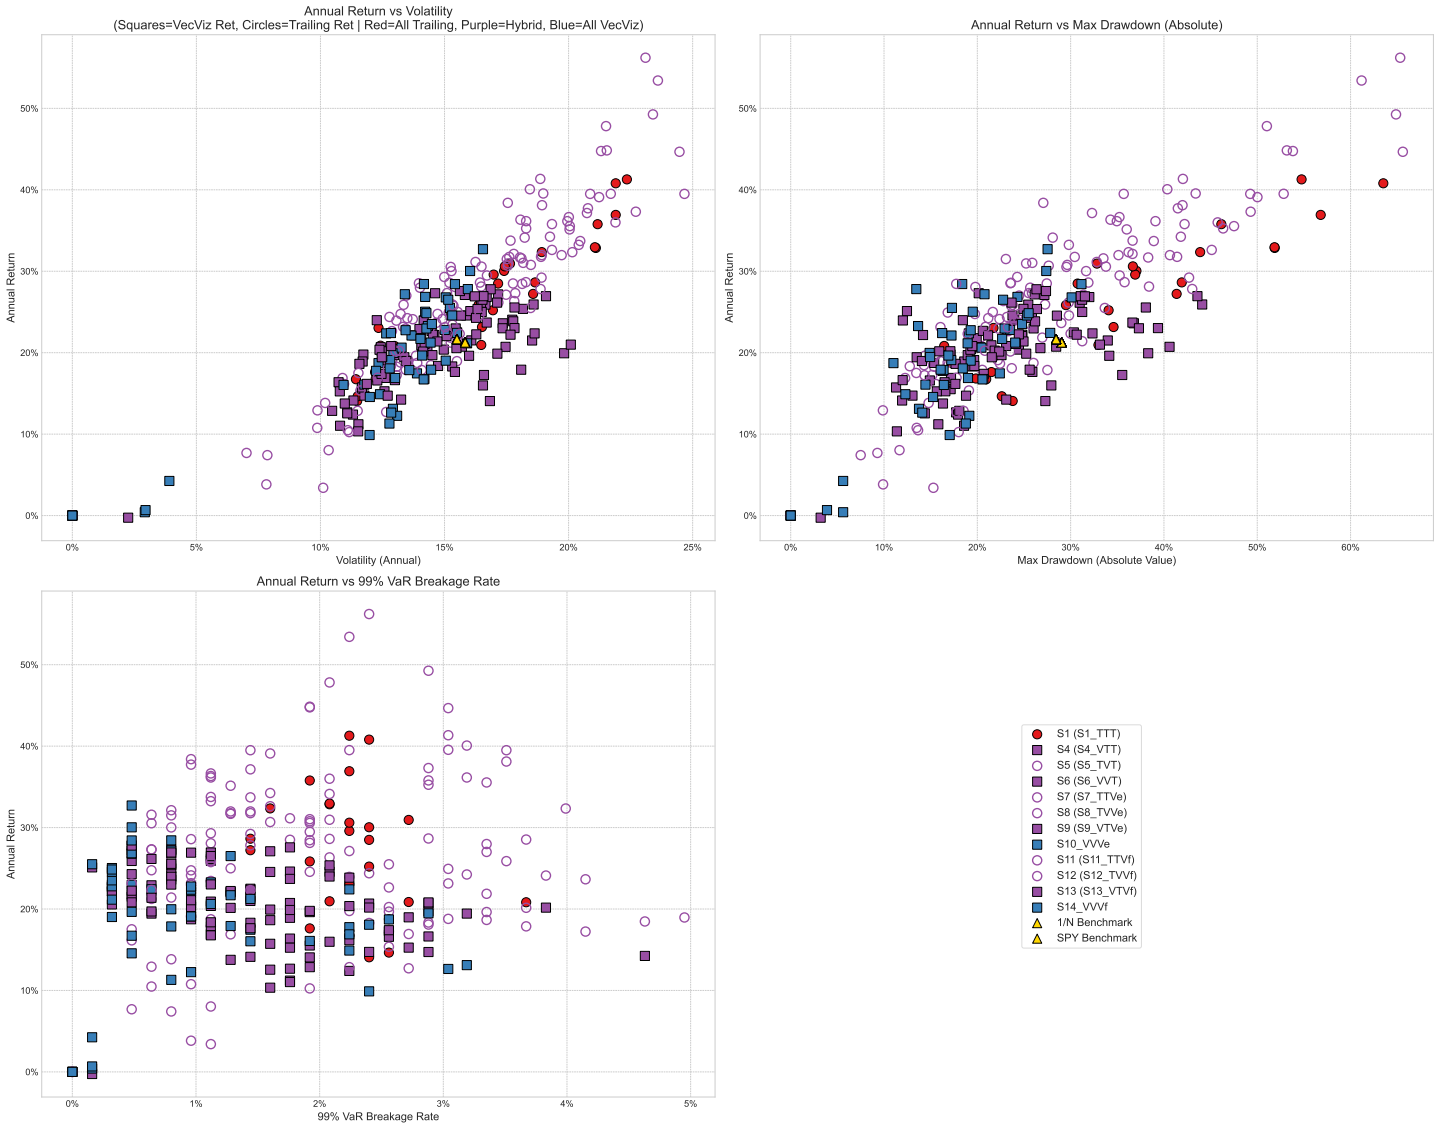

Appendix: Annualized Return (included in SummaryZ)

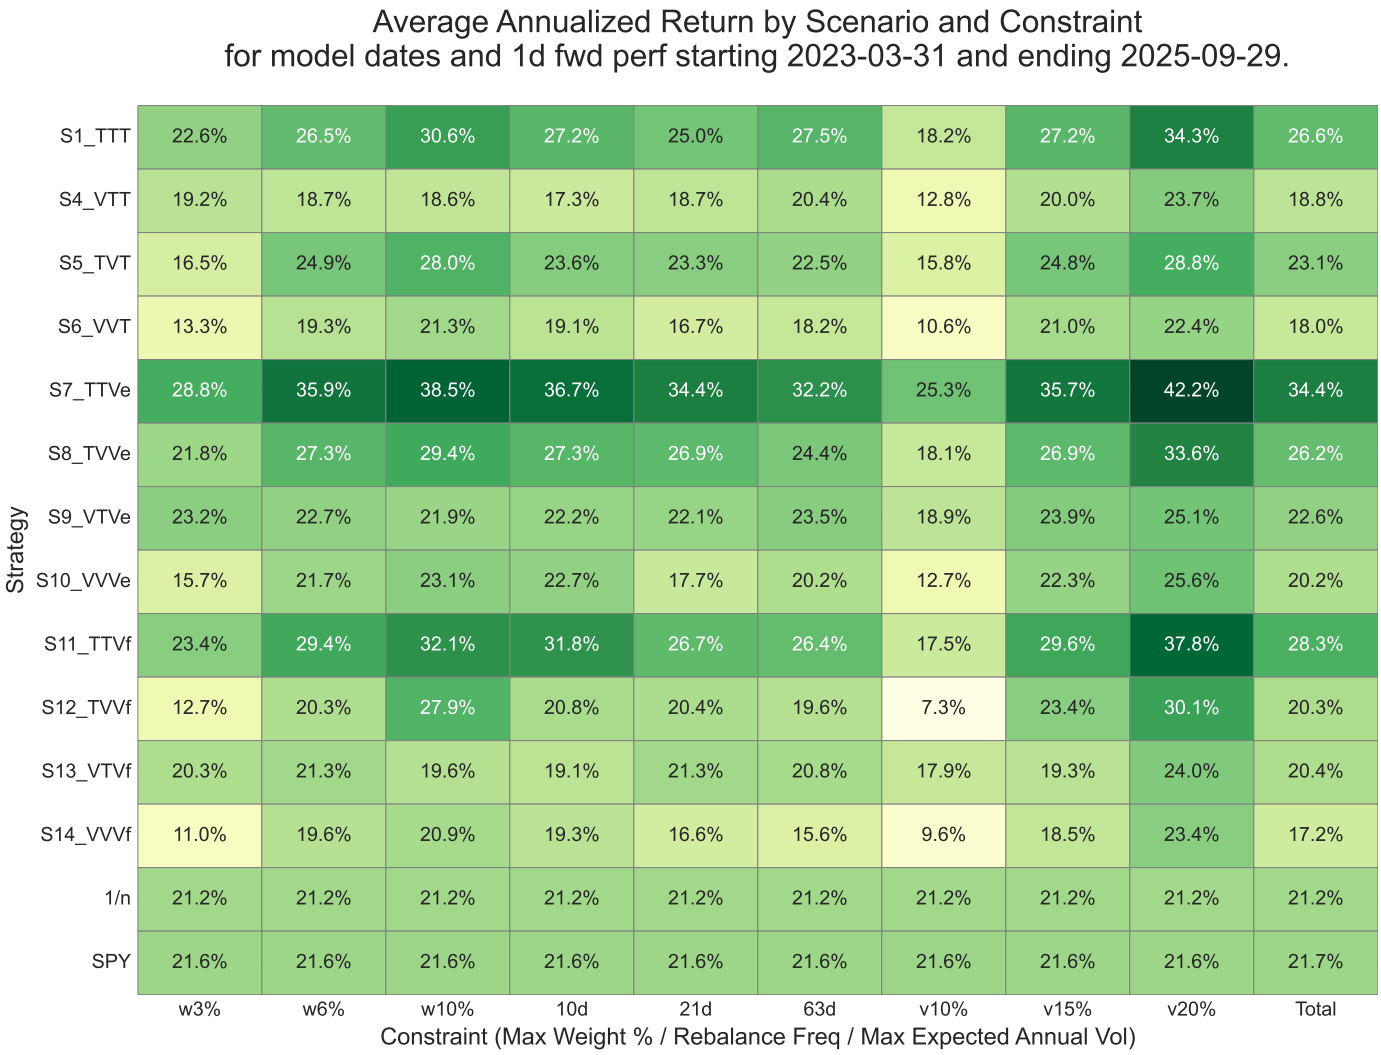

Ann_Return, is the annualized average daily strategy return. Returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Sharpe Ratio (includded in SummaryZ)

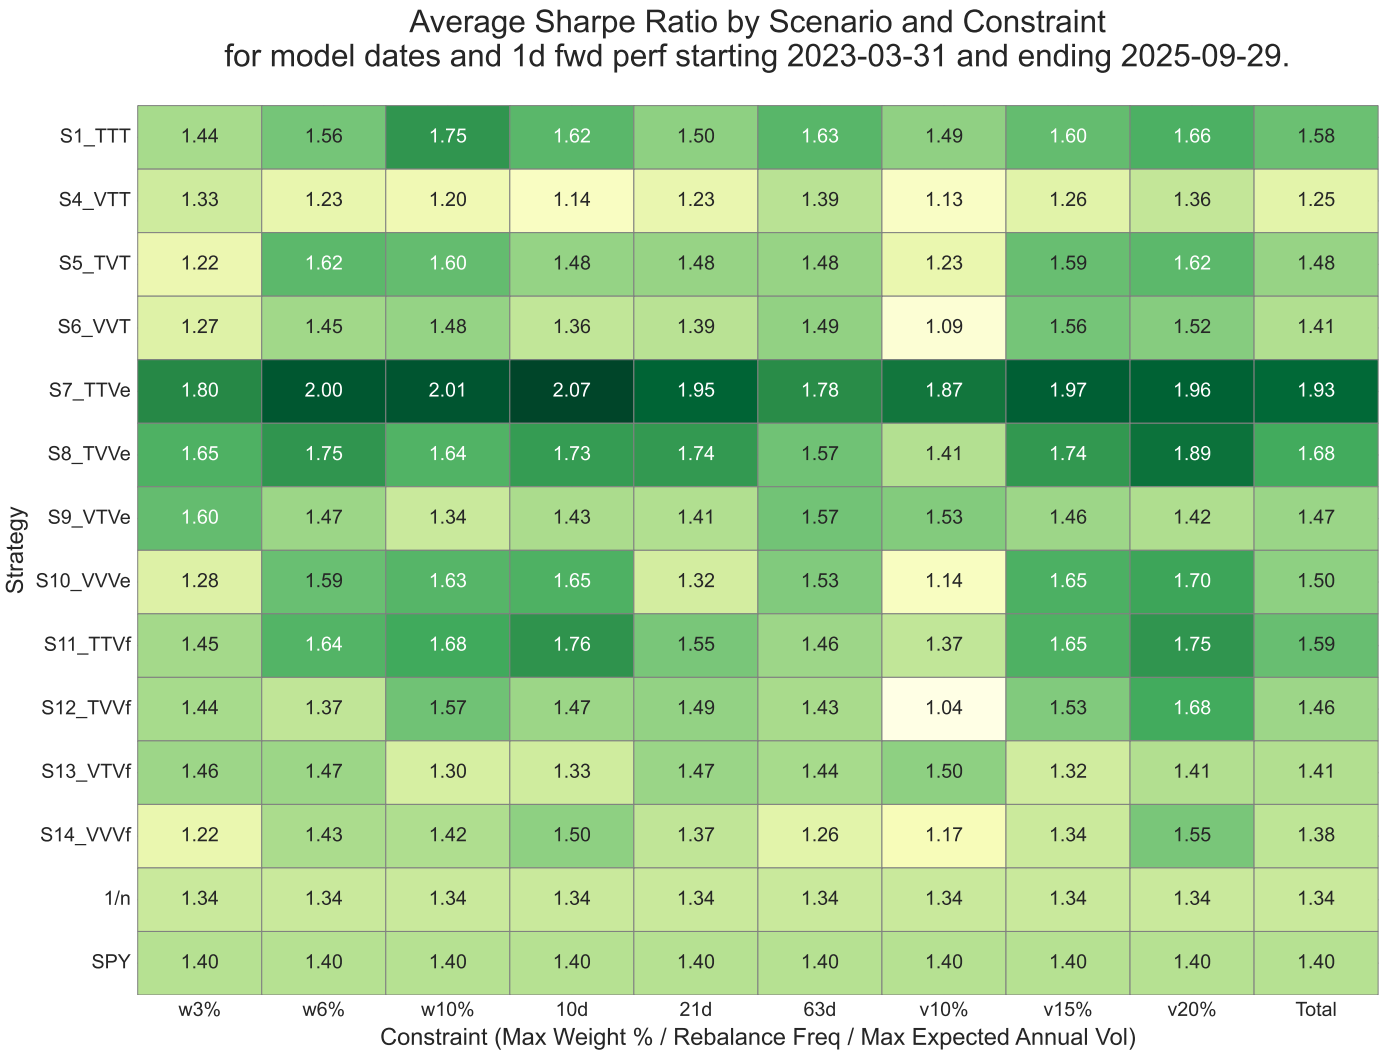

Sharpe Ratio is the annualized daily average strategy return / standard deviation of daily strategy returns. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Max Draw Down (included in SummaryZ)

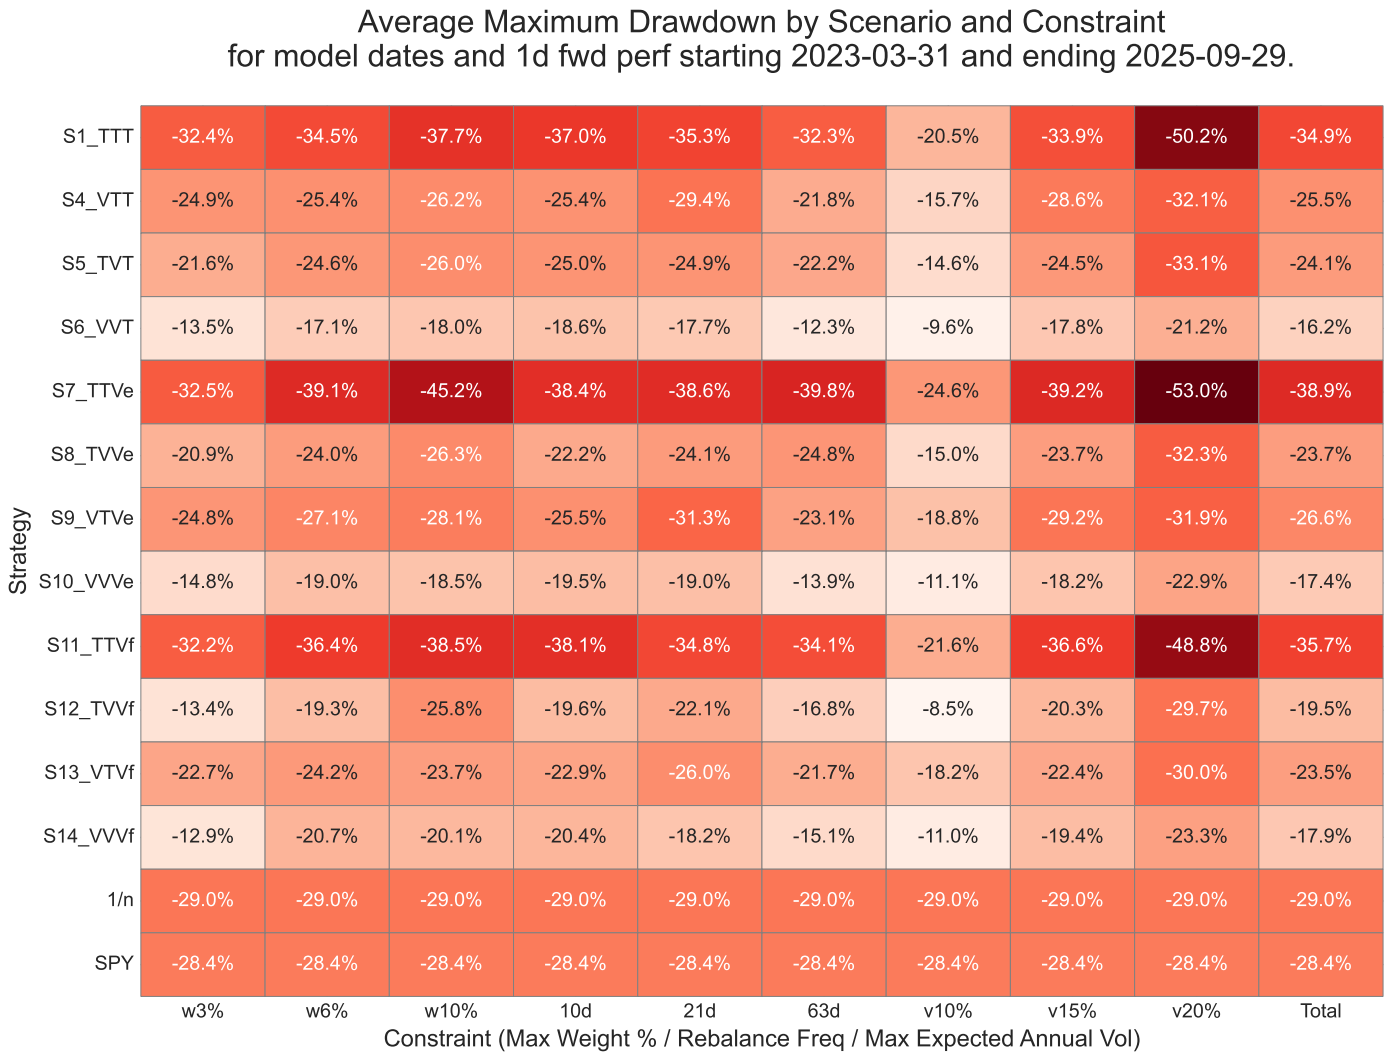

Max Draw Down (MaxDD), is the greatest decline from a peak in cumulative strategy return to a nadir in cumulative strategy return. Returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Calmar Ratio (included in SummaryZ)

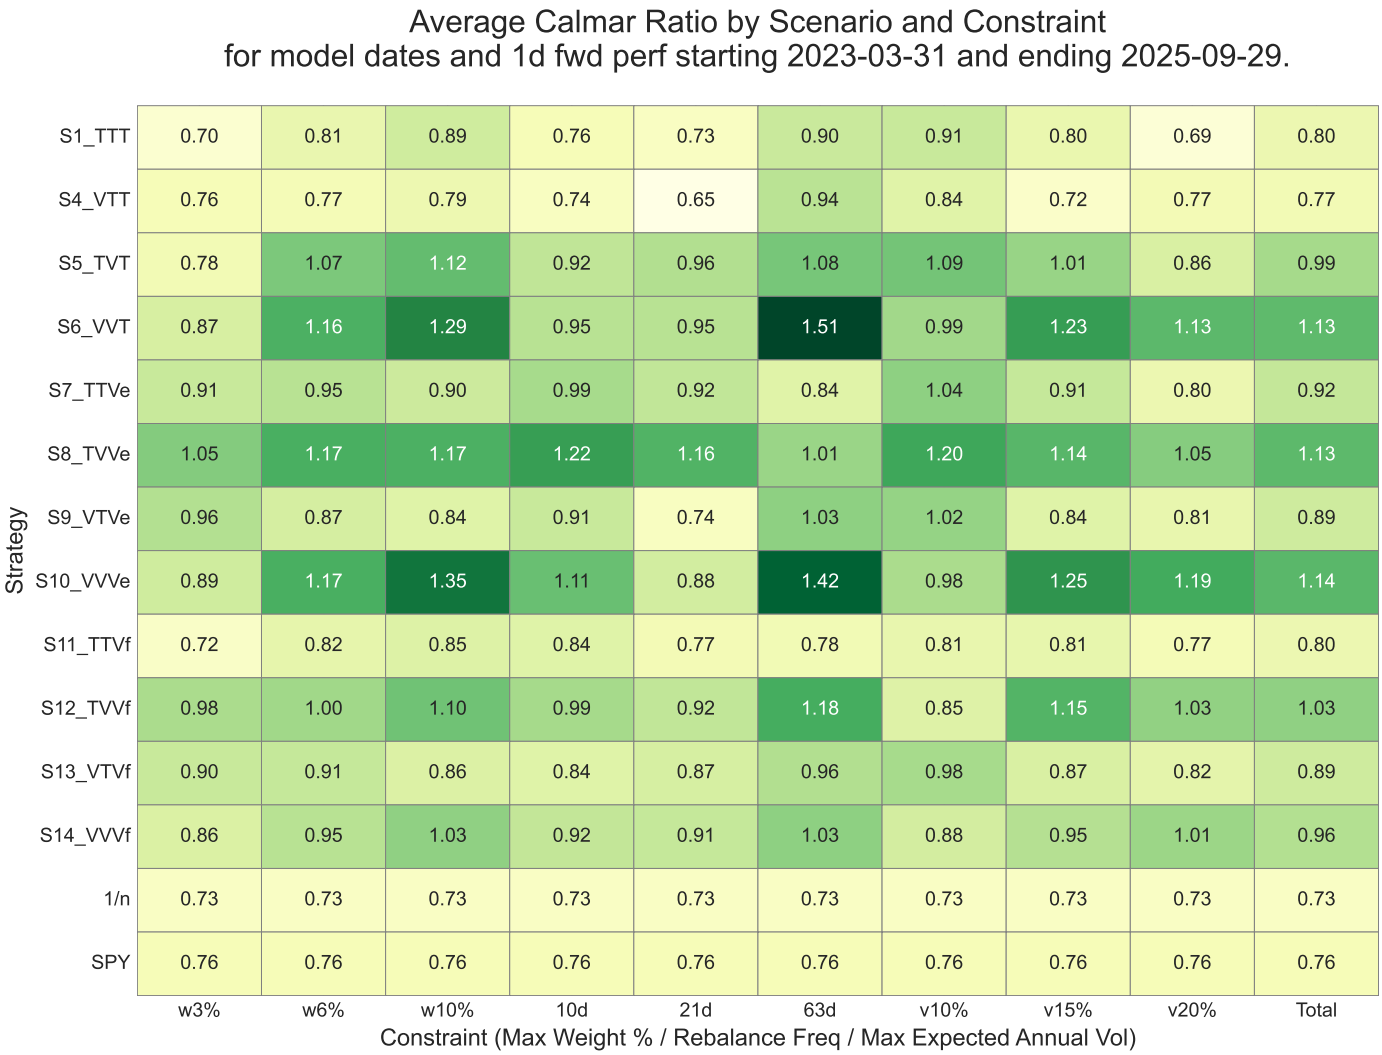

Calmar Ratio is the annualized daily average strategy return / Max Drawdown of cumulative daily strategy returns. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Multi-Factor Annualized Alpha (included in SummaryZ)

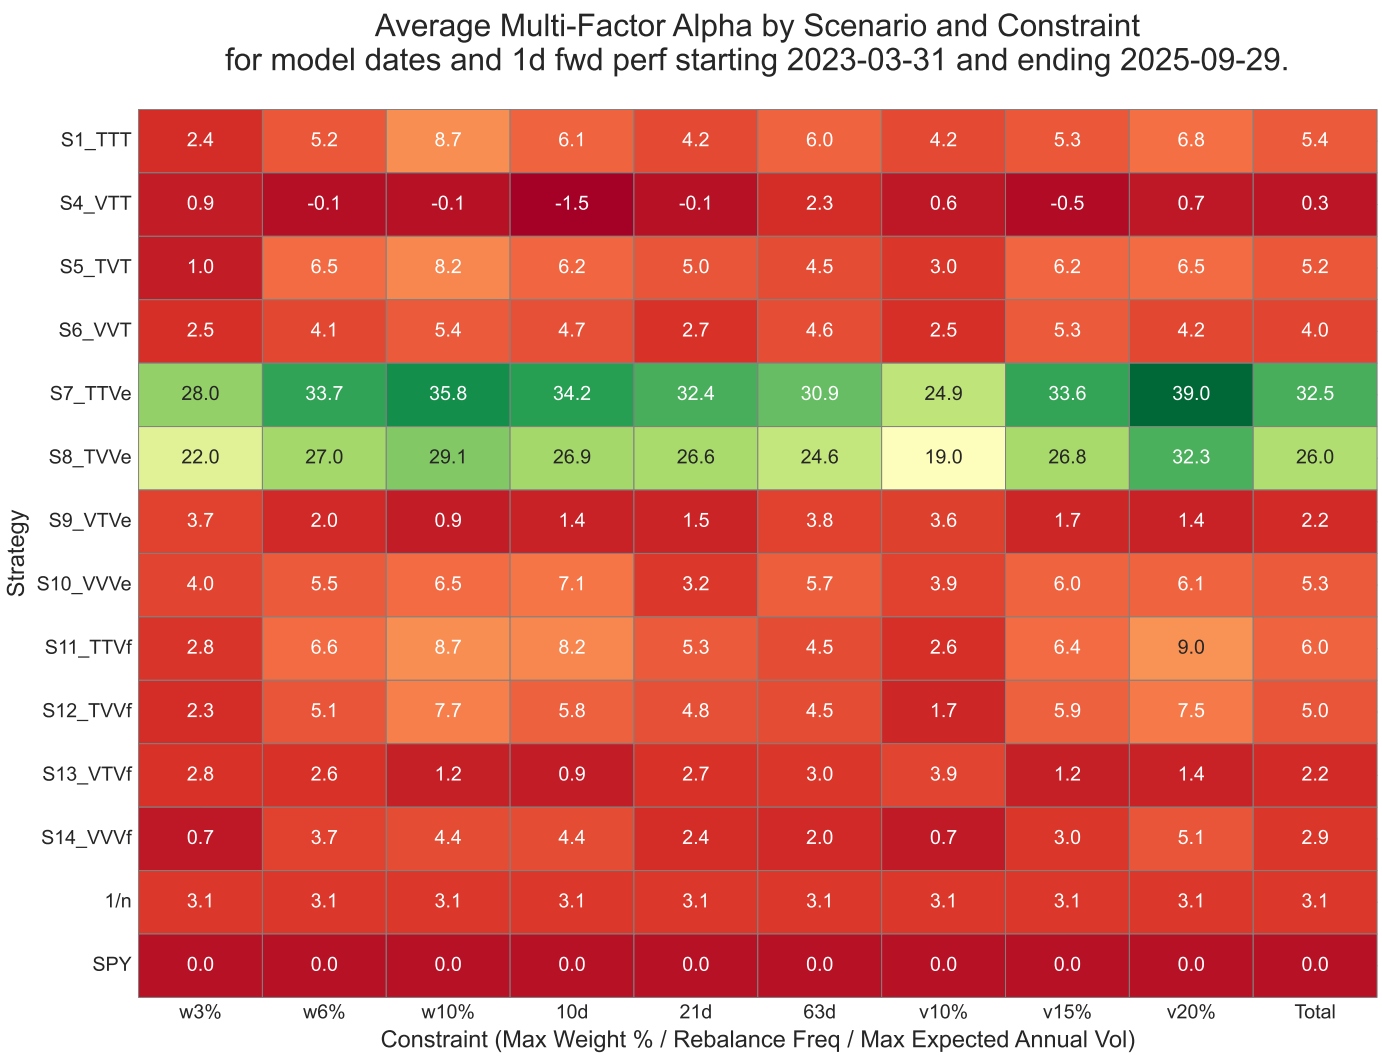

Multi-Factor Alpha is the annualized value of the intercept of daily strategy returns regressed upon the corresponding daily returns of SPY, MTUM, VLUE. Returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: 99% VaR Kupiec Test P-Value (included in SummaryZ)

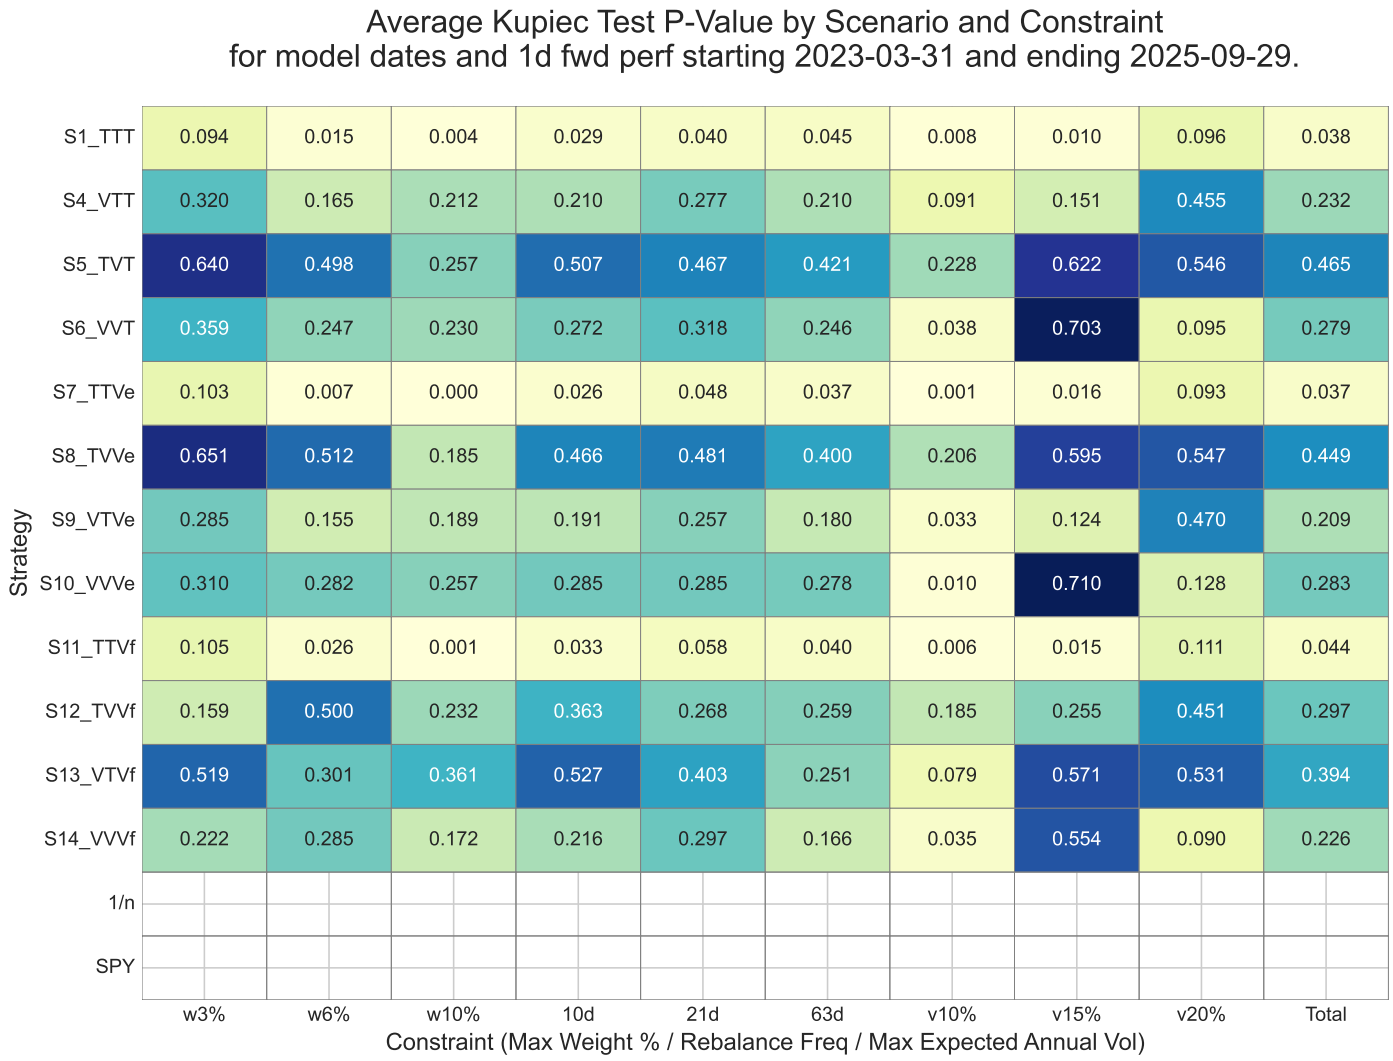

The P-Value of the Kupiec Proportion of Failures test statistic represents the probability that the null hypothesis that the 99% VaR breakage rate is consistent with expectations is true (higher Kupiec P-Value is better). ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Annualized Volatility

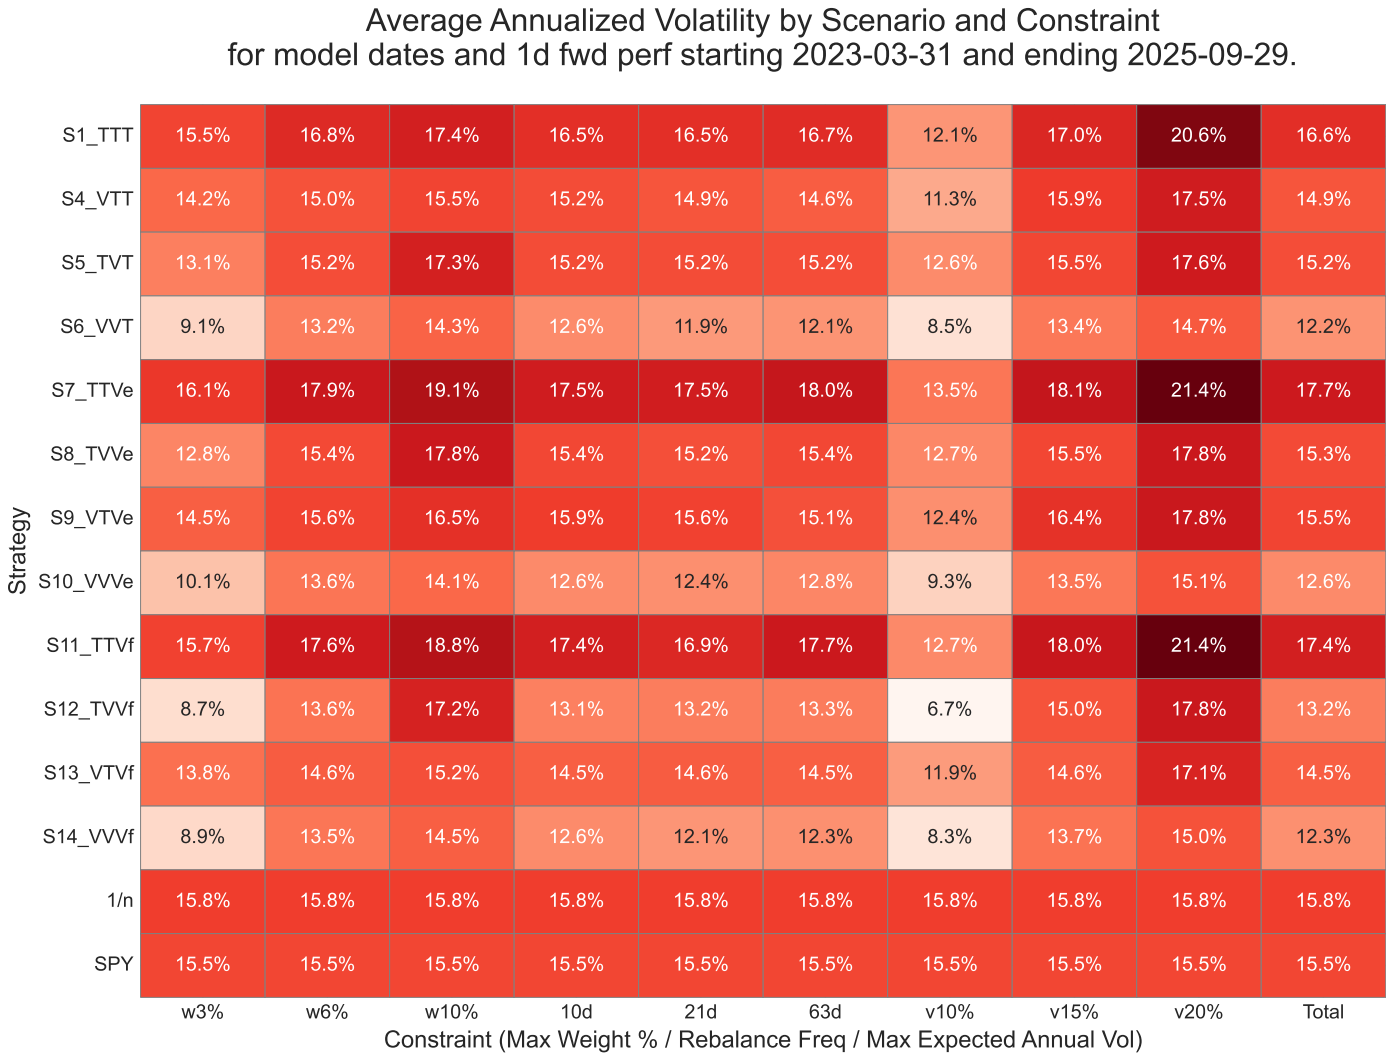

Vol, is the annualized standard deviation of daily strategy return. Returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. Annualized bolatility is Not included in SummaryZ because it is well represented in Sharpe Ratio, Max Draw Down and Kupiec P-Value. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: 99% VaR Breakage

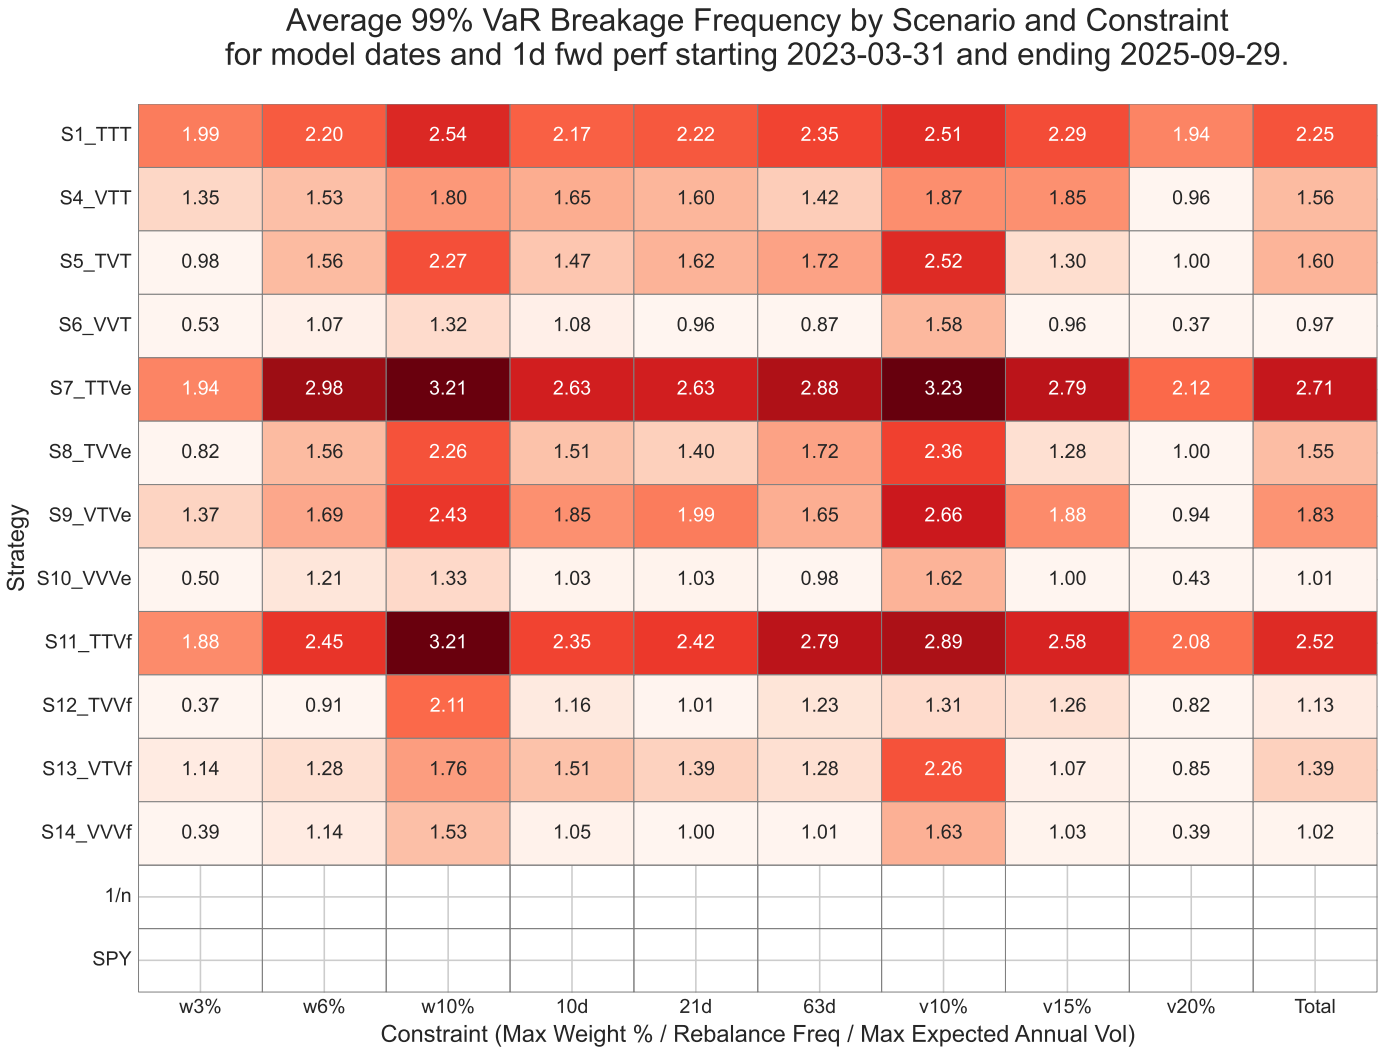

*### 99% VaR Breakage for a strategy represents the percentage frequency with which daily returns are less than -2.326*annualized Volatiltiy constraint / squrt(252). The ideal outcome would be 1.0%. Strategy returns on day d are based on weights determined on day d-1 and associated positions with initial cost at the closing value on d-1, using info avaliable up through day d-2, with no adjustment for transaction cost. 99% VaR Breakage is not included in SummaryZ becasue it is well covered by Kupiec Test P-Value. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.*

Appendix: Christoferson P-Value

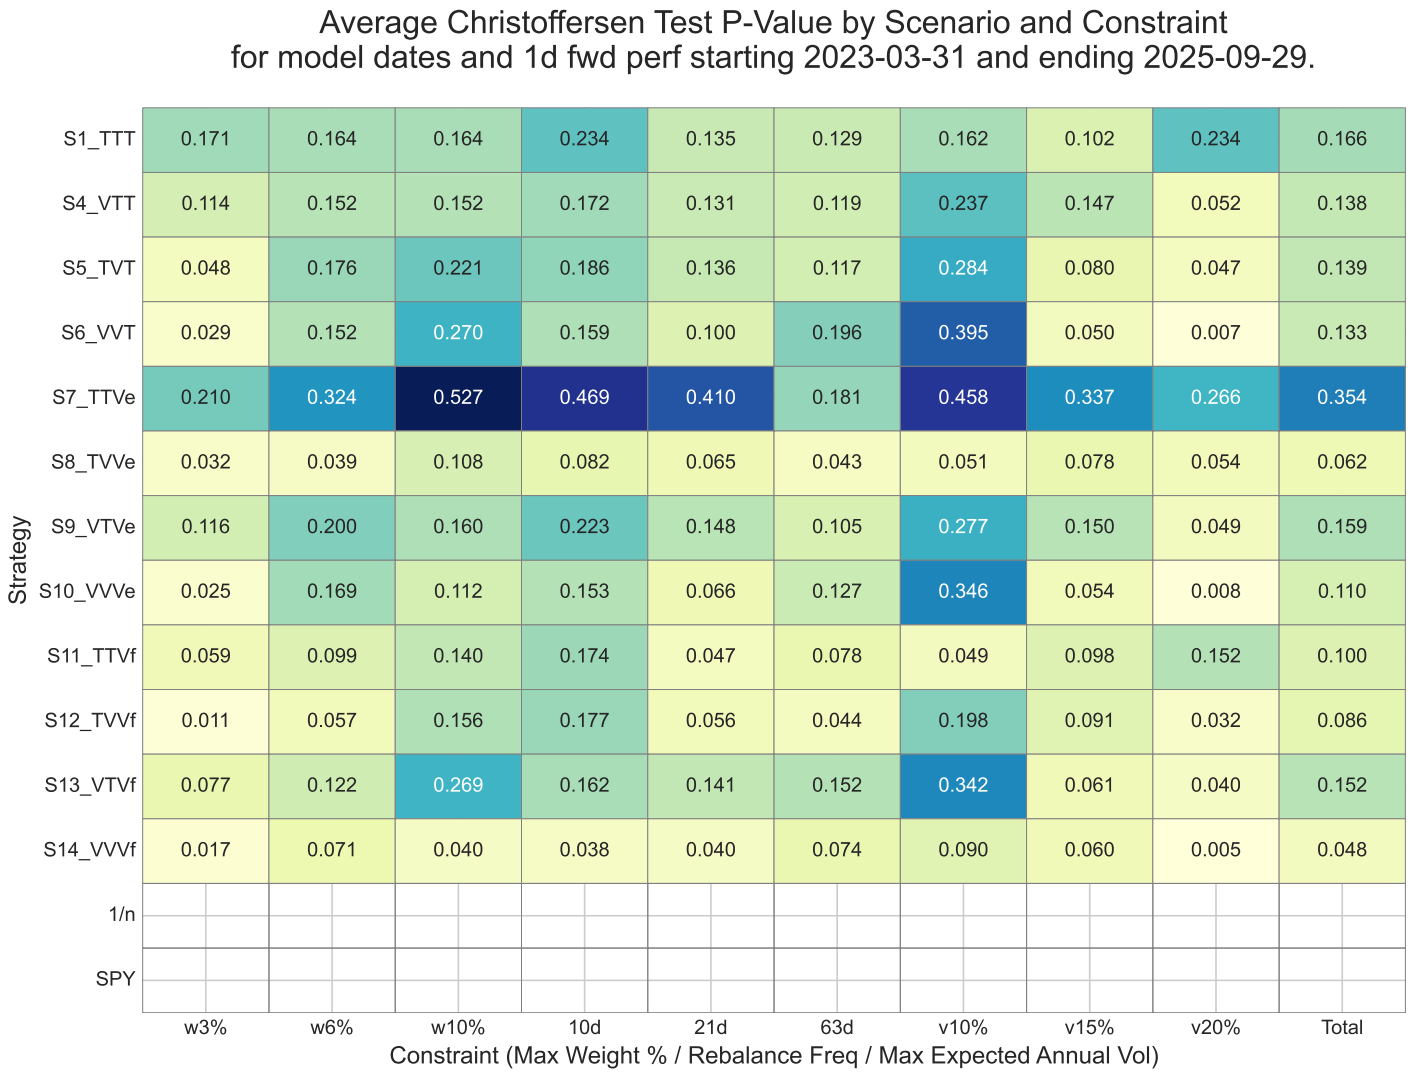

The P-Value of the Christoferson VaR Violation Independence test statistic represents the probability that the null hypothesis that the VaR model violations are independent is true (higher Christoferson P-Value is better). Christoferson is not included in SummaryZ, as its information is well represented by Max DrawDown. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Realized Vol MAE to Vol Constraint

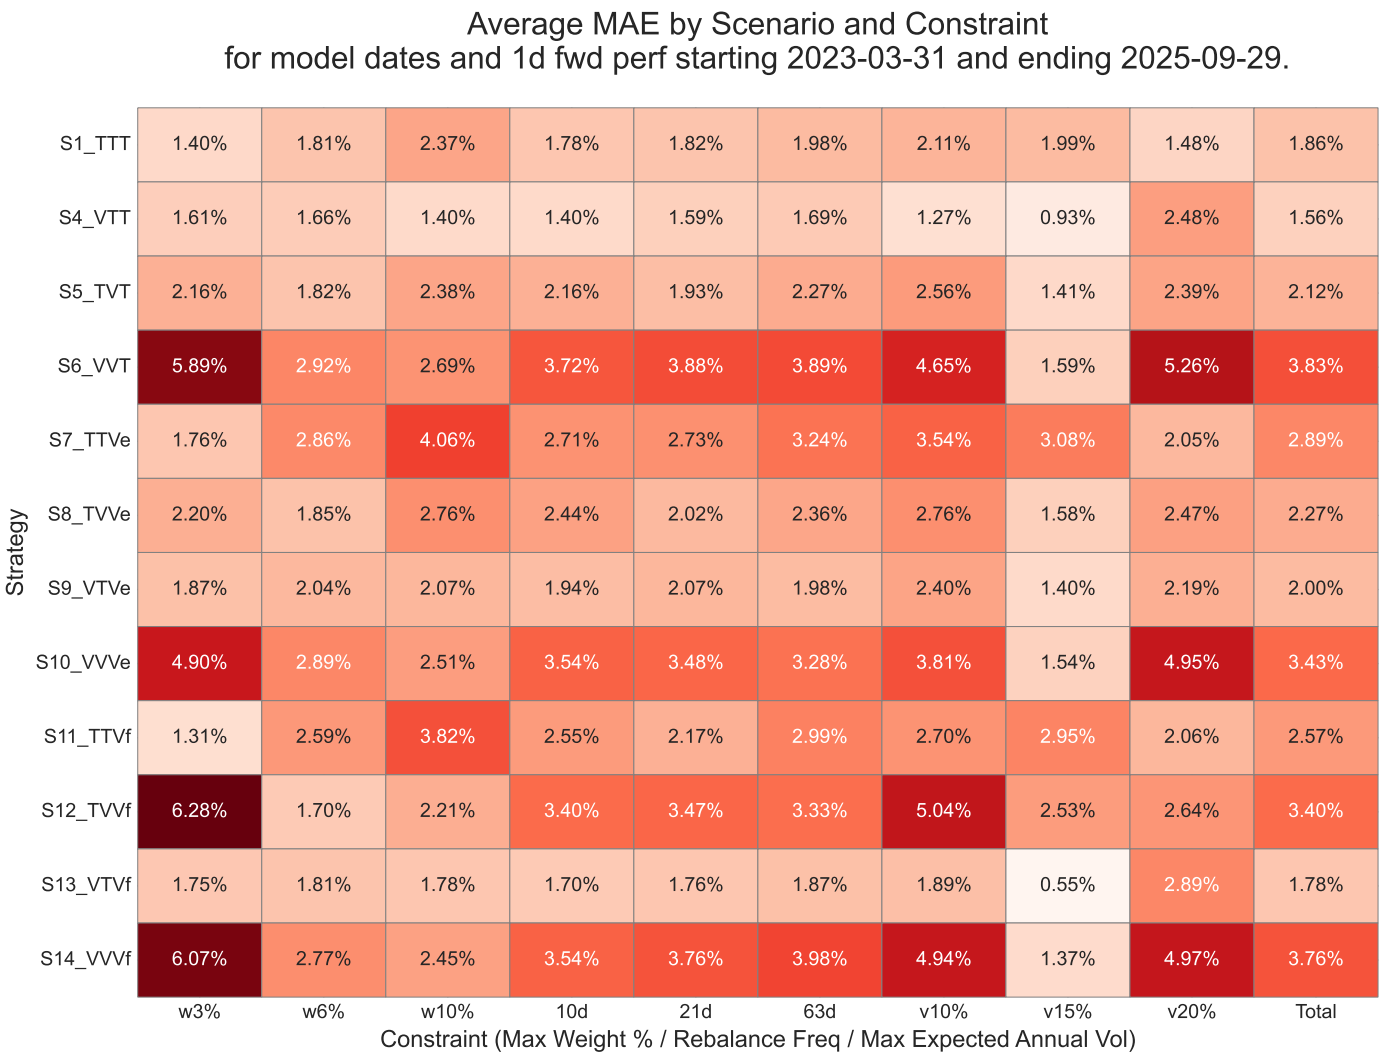

The Mean Absolute Error (MAE) relative to Vol Constraint is the MAE of realized strategy vol to the vol constraint (lower is better). Realized Vol MAE is not incuded in SummaryZ. It provides similar information as Kupiec P-Value. ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Rebalance Date 2-Way Turnover

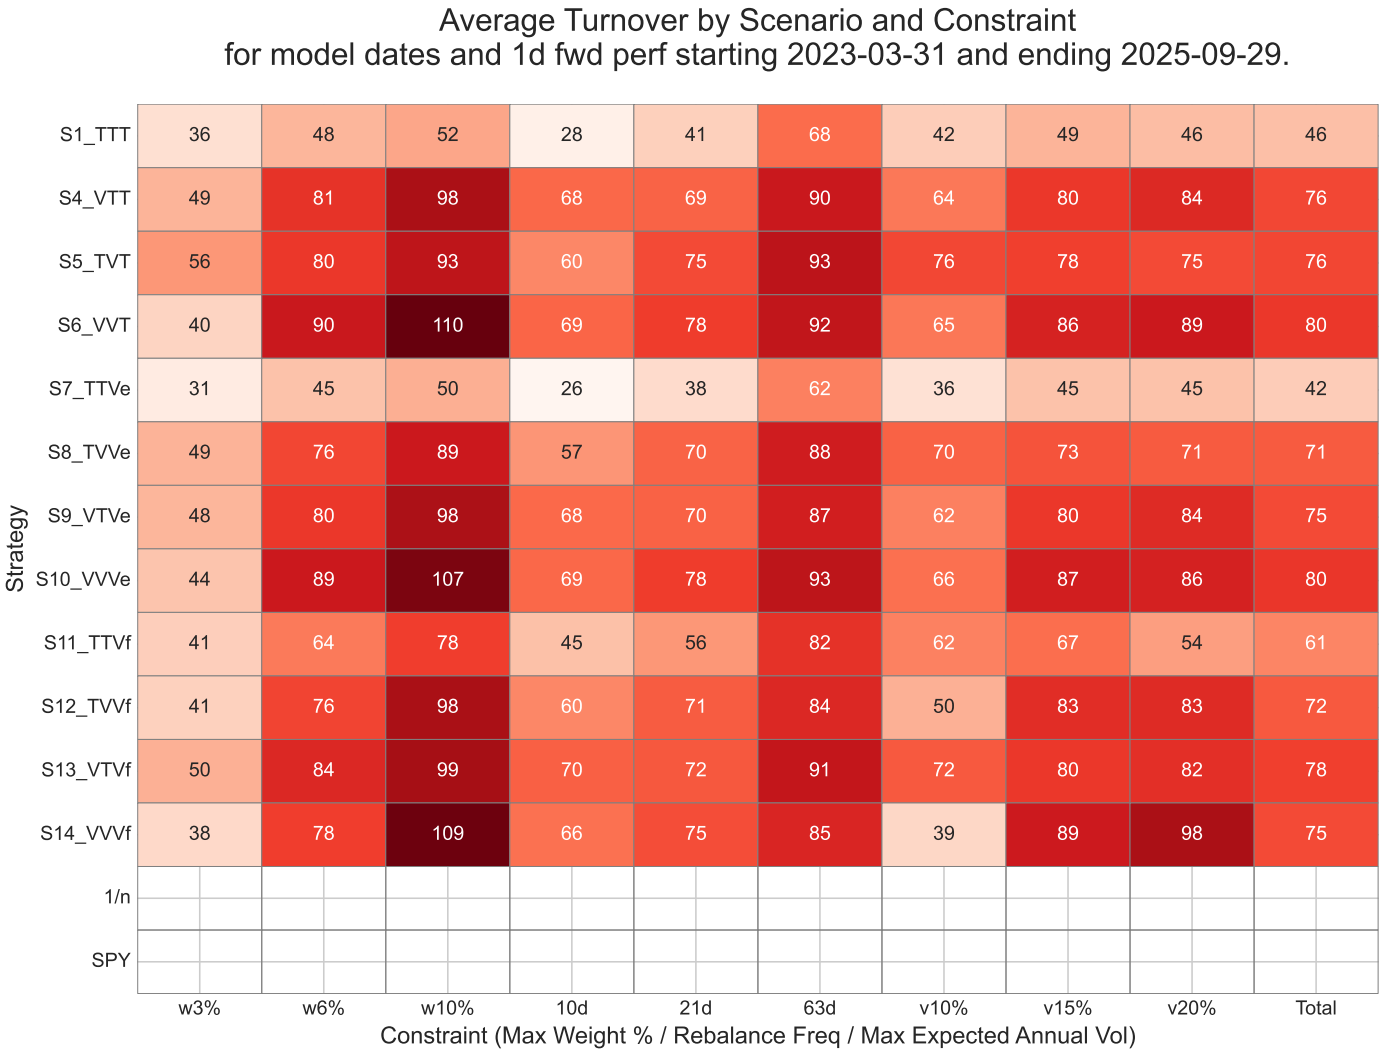

Turnover is the percentage of the portfolio the strategy buys and sells each rebalance date. For example, if 25 percent of the portfolio was bought and sold on the average reblance date then the turnover would be 50 percent. Note that these turnover statistics exclude the turnover associated with daily rebalancing to the most previously furnished set of optimized ticker weights between scheduled rebalance dates that is implicit in the return calculations for each strategy presented, including "1/n". ### STRATEGY rows are defined by input sources for return, volatility, and correlation: T = Trailing 252d; V = VecViz. VecViz correlation is either "e" (VecEvent-based) or "f" (VecViz analytic fingerprint). ### CONSTRAINT LABELS: those preceded by a "w" are max weight constraints, those followed by a "d" are rebalance frequencies, and those preceded by a "v" are max annualized expected portfolio volatility constraints. Cells in a given constraint column represent the average outcome for that constraint, across the 9 combinations of the other two constraints and rebalance frequencies. Cells in the "Total" column represent the average value across all 27 constraint outcomes for each strategy.

Appendix: Top Ticker Exposure Analysis

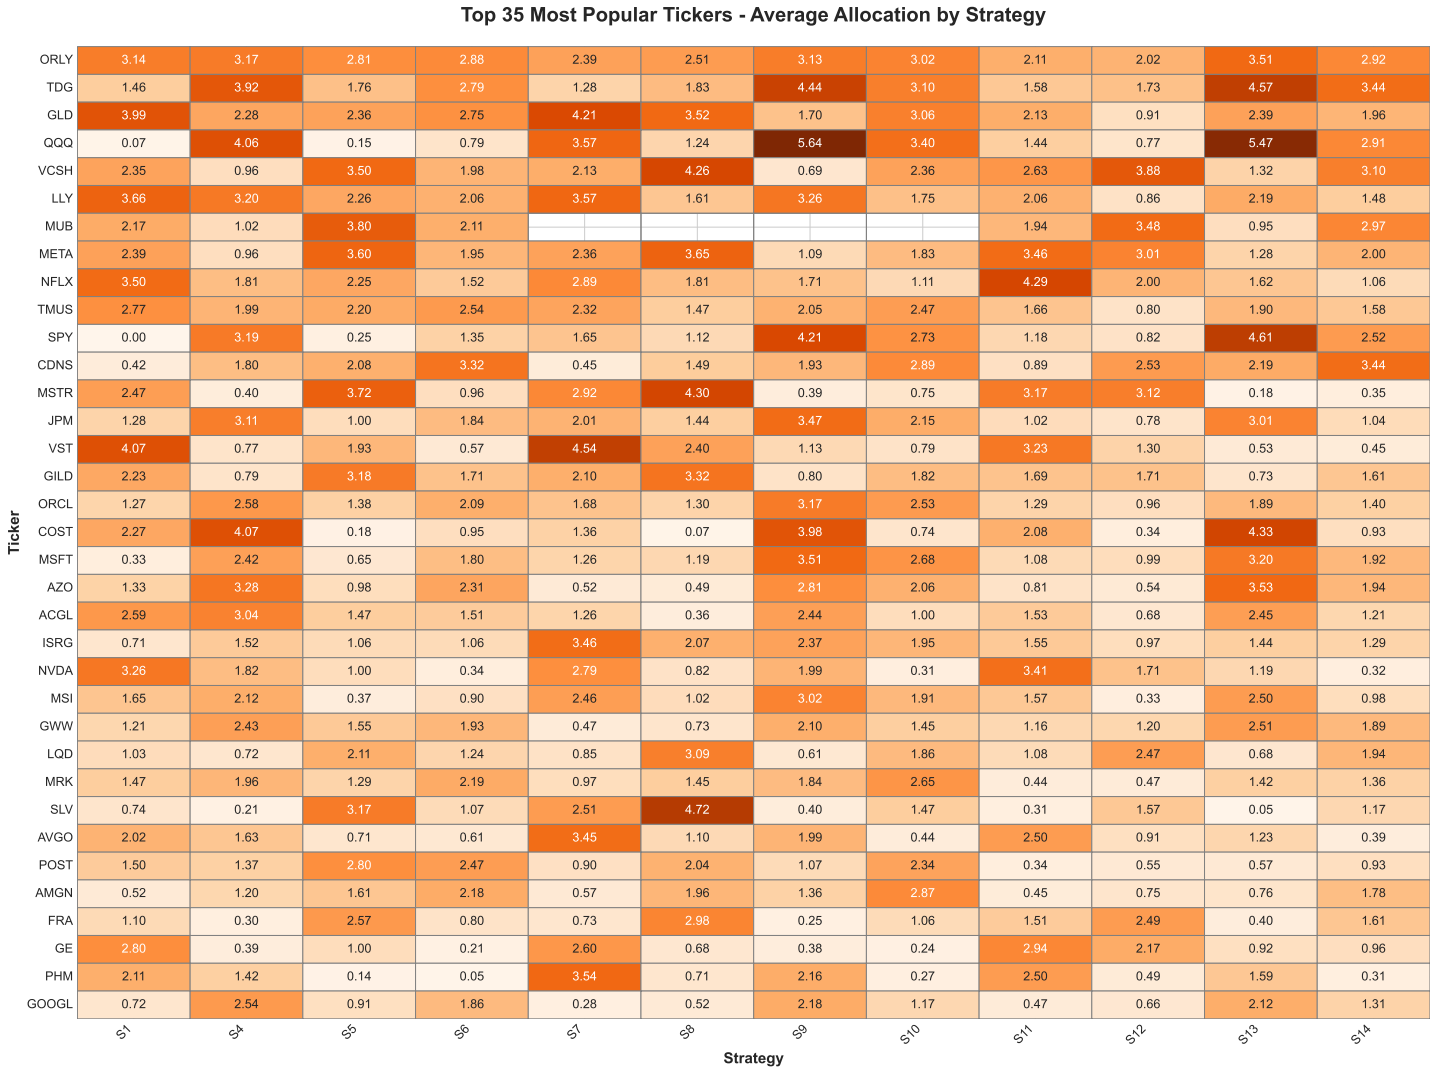

Appendix: Bottom Ticker Exposure Analysis

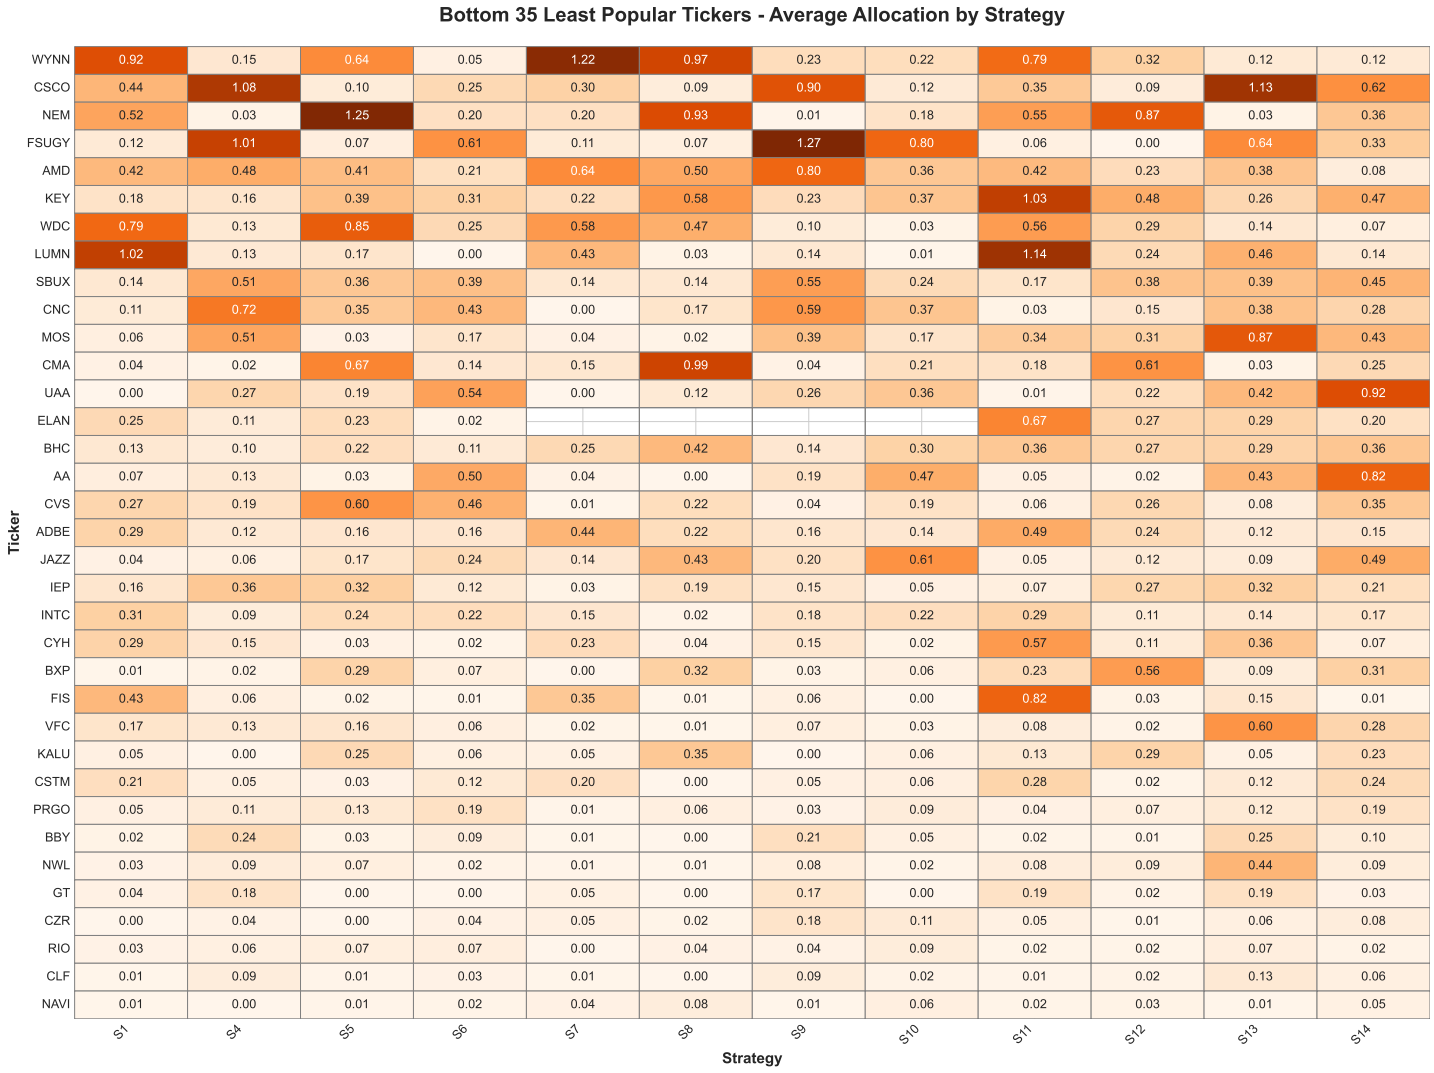

Appendix: Feature Utilization Across Regimes

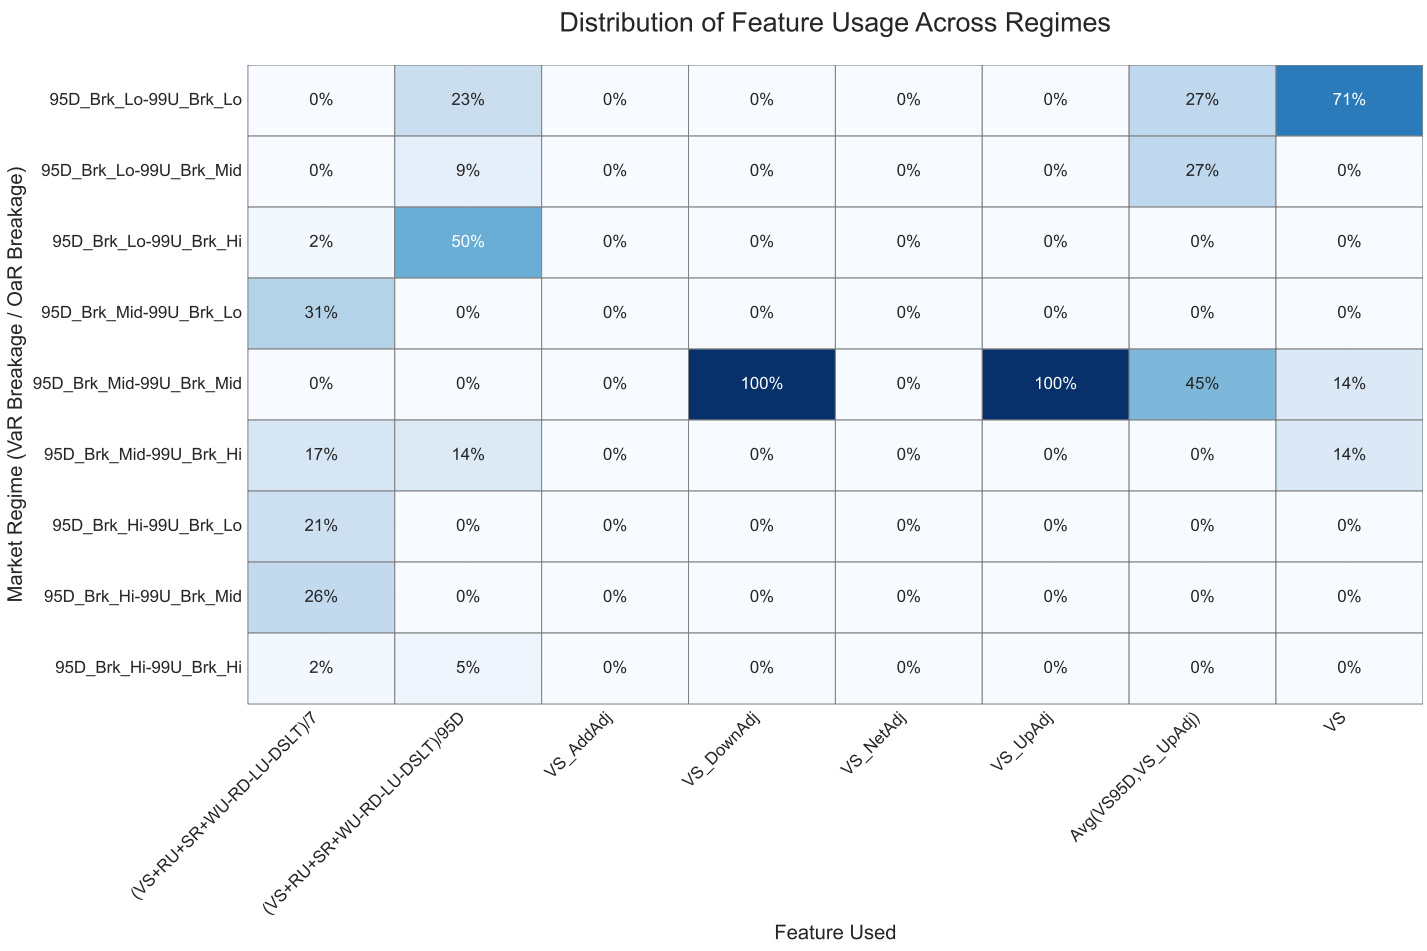

Appendix: Feature Utilization By Regime