Building a Diversification Quality Framework, Considering Dual-Dimension of Return and Volatility in Portfolio Construction

-

Jules-Grégoire Lafond

Republished for Kryptospshere®

This work is a republished version of the author's Master Thesis titled "**Building a Diversification Quality Framework, Considering Dual-Dimension of Return and Volatility in Portfolio Construction**" completed at NEOMA BS under the supervision of Maria Ruiz, CFA

Declaration of Academic Integrity

I, undersigned, M. Lafond Jules-Grégoire, certify that the data and information in the paper entitled *Building a Diversification Quality Framework, Considering Dual-Dimension of Return and Volatility in Portfolio Construction*. have not been plagiarized.

AI Declaration

In the context of this paper backtested portfolio strategies were done which needed the use of large databases handled by using Python programming code.
The backtesting program has been entirely coded using AI assistant under the implementation guidelines given by author of the paper
The code is available upon request at julesgregoire.lafond.22@neoma-bs.com

This work is a republished version of the author's Master Thesis titled "**Building a Diversification Quality Framework, Considering Dual-Dimension of Return and Volatility in Portfolio Construction**" completed at NEOMA BS under the supervision of Maria Ruiz, CFA

Table Of Contents :

Abstract

I) Conceptual Foundations
- A) Hypothesis Introduction
- B) Peer Review
- C) Conclusion

II) Introducing new ratios to measure Risk
- A) Integrating Volatility Dispersion Index (VDI)
- B) Introducing Volatility Stability Ratio (VSR)
- C) Achieving Volatility Diversification Ratio (VDR)

III) Building the Diversification Quality Framework
- A) Diversification Dual Dimension Table
- B) Proposing a Fully Diversified Portfolio
- C) Nuancing the Framework

IV) Backtested Portfolios
- A) Dataset Treatment & Processing
- B) Signal Normalisation
- C) Convex Optimisation and Portfolio Formulations
- D) Backtest Results and Performance Analysis

V) Conclusion

VI) References

VII) Appendices

This work is a republished version of the author's Master Thesis titled "**Building a Diversification Quality Framework, Considering Dual-Dimension of Return and Volatility in Portfolio Construction**" completed at NEOMA BS under the supervision of Maria Ruiz, CFA

Abstract :

This paper inspects how to improve Risk Factor exposures through the lens of volatility diversification. The idea is to consider Risk Factor exposures in portfolio construction via the measure of correlation of the assets' volatilities. The starting point is the following : When assets see their volatilities increase together, it highlights the shared sensitivity of those assets to the same underlying risk. Measuring this shared sensitivity is done through two new metrics "Volatility of Volatility" (VoV) which is used to compute the "Volatility Dispersion Index" (VDI) which assesses how much do assets' volatilities spike together in a portfolio. This pushes to review the definition of Systematic Risk and Diversifiable Risk and what this implies for portfolio construction as portfolios might not be as much diversified as they claim.

Then, to improve risk factor diversification and resilience this paper examines how to build an improved diversified portfolio thanks to a Dual-Dimension Diversification Framework based on correlation of returns and correlation of volatilities by integrating these new metrics, namely Volatility Dispersion Index (VDI) and Volatility of Volatility (VoV) in the Diversification Ratio (DR) formula, developed by Yves Choueifaty & Yves Coignard (2008) Towards maximum diversification [1] to produce new measures for measuring volatility stability and volatility diversification in portfolios. The VoV replaces the weighted average volatility component and the VDI replaces portfolio volatility component.
This paper then evaluates the performance of portfolios constructed around these metrics, Volatility Diversification Ratio (VDR) and a DR-VDR portfolio, which tries to maximise both, compared to portfolios constructed using the traditional Diversification Ratio methodology.

Keywords :

Risk Factor investing, Diversification, Volatility of Volatility (VoV), Volatility Correlations, Systematic & Diversifiable Risk

This work is a republished version of the author's Master Thesis titled "**Building a Diversification Quality Framework, Considering Dual-Dimension of Return and Volatility in Portfolio Construction**" completed at NEOMA BS under the supervision of Maria Ruiz, CFA

I) Conceptual Foundations

A) Hypothesis Introduction

The pursuit of optimal portfolio construction has long been guided by the principles of Modern Portfolio Theory (MPT) of 1952, where the diversification of assets serves as the fundamental pillar for mitigating risk. While traditional metrics have relied on historical data to quantify risk and correlation, the increase in sophistication of financial markets and the prominence of derivative instruments have led an era where predictability and stability in investing has drastically become a priority for major financial institutions to plan cash flows and returns over long periods. Among these, the Diversification Ratio (DR) stands as a gauge of a portfolio's efficiency in diversifying its assets and reducing portfolio concentration. It is able to measure the reduction in volatility achieved by combining imperfectly correlated assets. Conventionally calculated using correlation of returns for its denominator, the DR provides an insight of return diversification efficacy of a portfolio.

However when talking properly about risk itself, DR does not really measure how well a portfolio is diversified from risk factors themselves, as it studies correlation of returns, not correlation of volatilities.

This is why, this paper investigates several pivotal modifications to this established metric:

- **First,** the substitution of portfolio volatility (σ_p) with volatility of aggregate volatility, or rather **Volatility Dispersion Index (VDI)** which is a metric specially built to analyse how volatilities of the constituents of a portfolio move together, in the denominator of the Diversification Ratio. Which will give a new tool to measure volatility stability in a portfolio, **Volatility Stability Ratio (VSR),** which measures a portfolio's volatility stability or in other terms, VSR measures how stable and predictable the portfolio's overall risk level is, by comparing the average volatility of the assets to the instability of that average.
- **Second,** the substitution of weighted average volatility of assets by the weighted average volatility of volatilities (WVoV) which will give a ratio able to measure if the portfolio's total stability risk is driven by the asset's individual instability or by their coordinated co-movements when getting volatile. In other terms the **Volatility Diversification Ratio (VDR)** measures if the portfolio's stability (the portfolio's volatility of volatility) is mainly driven by diversifiable risk or by systematic risk. Thus this ratio measures the diversification to volatility of risk of each asset.

These approaches alter the interpretation of the ratio, where Diversification Ratio measures the diversification of returns, these ratios measure the stability of risk for the average risk taken in the portfolio (VSR) and the nature of instability, Systematic Risk or Diversifiable Risk, present in the portfolio (VDR).

This work remains an introduction to the subject and analyses the surface of this modification, The analysis of this modified Diversification Ratio is assessed with a backtest across universes of 500 Stocks respectively through 13 years of market data. The following section presents the foundational concepts, and reviews the foundational work and more modern literature about risk. Followed by the new ratios' properties and methodology before testing portfolio strategies using these tools and delving into a comprehensive analysis of the results, where the impact of this seemingly simple substitution is quantified. And conclusions are drawn from it regarding risk management and strategic trading.

B) Peer Review

a. Foundational Theories

Risk is one of those concepts that finance has never quite finished defining. What is striking about the literature is not the technical progression, each model more sophisticated than the last, but the degree to which each redefinition quietly shifted what the word was supposed to mean in the first place, often without acknowledging it.

Markowitz (1952) [2] made the foundational move: risk is variance, and variance is not a property of individual assets but an emergent property of how assets relate to one another. This was genuinely radical, it transformed diversification from folk wisdom into a mathematical object and made portfolio construction an optimisation problem. The limitation, obvious in retrospect, is that Markowitz treats all variance as equal. A portfolio that oscillates unpredictably around a high mean looks identical in his framework to one that oscillates around a low mean with the same amplitude. Sharpe (1964) [3] attacked this directly: not all variance is economically equivalent, because diversification can eliminate the idiosyncratic component for free. What remains, systematic co-movement with the market, captured by

beta, is the only variance that should command compensation in equilibrium. This was conceptually clarifying but introduced a different problem: it reduced the entire complexity of systematic risk to a single number.

Black, Scholes and Merton (1973) [4][5] sidestepped the cross-sectional question entirely and approached risk from a continuous-time pricing perspective, treating volatility not as a portfolio characteristic but as a directly priceable quantity embedded in the dynamics of the underlying process. The intellectual move was elegant, implied volatility became the market's own forward-looking risk estimate, observable in real time. But the framework's load-bearing assumption, that sigma is constant, was also its most fragile. Markets had known for decades that volatility clusters; BSM essentially required ignoring that to make the mathematics tractable. Engle (1982) [6] formalised what practitioners already suspected: conditional variance is not constant but endogenously evolving, large shocks beget large shocks, and any model that treats volatility as static is fundamentally misspecified. ARCH was important not just as a technical contribution but as a conceptual one, it established that risk itself requires its own dedicated modelling apparatus, separate from the modelling of returns.

Fama and French (1992) [7] then returned to the cross-sectional question Sharpe had claimed to close, and found it very much open. Market beta, the sole priced risk factor in CAPM, produces a statistically flat relation with average returns in the data. Systematic risk is not one-dimensional but a vector, size, financial distress, and market sensitivity each contribute independently. This finding is the intellectual foundation of factor investing, but it also carries a less comfortable implication that the literature has been slower to absorb: if the relevant risk factors are empirically discovered rather than theoretically derived, there is no principled stopping point. The factor zoo that followed, momentum, quality, low volatility, profitability, is the direct descendant of Fama and French's empirical falsification of CAPM, and it raises the question of whether systematic risk has a stable structure at all, or whether it is simply whatever the cross-section happens to price at a given moment.

What this progression reveals, taken as a whole, is that each successive framework did not so much solve the problem of risk measurement as relocate it. Markowitz moved risk from the asset to the portfolio. Sharpe moved it from total variance to systematic variance. BSM moved it from the portfolio to the price process. Engle moved it from a parameter to a state variable. Fama and French moved it from one dimension to many. The question this paper takes up, whether the volatility of volatility constitutes a genuinely distinct, independently structured dimension of risk, sits naturally in this lineage. It is not a departure from the tradition but its next logical step.

b. Modern Literature review

Maximum Diversification Ratio is a ratio defined by the weighted average volatility of the assets composing a portfolio divided by the volatility of the portfolio.

$$Diversification\ Ratio\ (DR) = \frac{\sum_{i=1}^{n} w_i\, \sigma_i}{\sigma_p}$$

This ratio measures the diversification of returns amidst a portfolio of assets based on the correlations and volatilities of the assets. The higher the ratio, the higher the diversification, meaning that the portfolio is able to make each asset's weight contribute equally to the portfolio's volatility of return. 20% of stock X has the same impact than 2% of stock Y for instance. The ratio measures risk associated to return, and is used as a risk reduction tool. Portfolios constructed using it tend to have a bias towards lower-volatility return and lower-correlation assets. The optimization process naturally allocates more weight to assets that contribute less to overall portfolio return risk and have a greater potential to offset the volatility of return of other assets. This can lead to portfolios that are inherently more defensive and may offer better downside protection during periods of market stress.

And as a matter of fact, this construction tends to have two main problems :

- By definition, DR portfolio will diversify in order to avoid concentrated positions, which leads to underperformance against the other during concentration phases in markets, when a few stocks outperform and carry the overall indices performance over time. With the most blatant example being the dominance of the technological sector over the past decade.

- It does not really study risk itself, but risk associated to return. Meaning that it does not really help reducing actual risk when market stress occurs, as correlation of assets converge towards 1 when it happens. The portfolio is well-diversified in calm markets but not when portfolio managers really needs diversification in times of crisis. Indeed assets can have the same risk

factor, which will be partly considered by DR given it also looks at correlation of returns among assets.

By construction, Maximum Diversification Portfolio can have high exposures to specific risk factors, exposing a portfolio to a surface diversification - diversification of return - but not diversification of risk.

This is precisely what Volatility Diversification Ratio (later introduced) tries to solve, addressing the mathematical and operational challenges of accurately decomposing and attributing portfolio risk to its underlying drivers which is the practical implementation of the layered framework that will be introduced later. This is done through the computation of Volatility-of-Volatility (VoV), which will be detailed later on in the calculation part of this paper. This represents a true issue since in modern literature, Rosen and Saunders (2010) [8] confront this challenge directly in the credit risk context, demonstrating that a portfolio's loss profile is a non-linear function of its underlying risk factors and that standard linear attribution techniques are therefore structurally inadequate. Their Hoeffding decomposition framework, which isolates individual factor contributions while capturing higher-order cross-factor interaction terms, provides the formal analytical precedent for the risk factor attribution logic underlying Figure c (see below). In the end the asset class to which this applies does not matter as all asset classes have different risk factors impacting them, meaning the framework and the indicators we'll introduce later on can also be used in other asset classes not only equity.

However, one could object saying though that to hedge or diversify against risk, volatility products such as the VIX could be used, so to have a leg exposed to the market's stress condition which would provide a protection or at least dilute portfolio risk if the market enters a high-volatility phase. However, Alexander and Korovilas (2011) [9] illustrate the operational hazards that emerge when volatility itself is treated as a diversifying instrument: the persistent contango structure of VIX futures imposes a continuous roll-cost drag that renders permanent volatility allocations net-negative over extended horizons, demonstrating that the dynamic, time-varying nature of factor sensitivities, acknowledged in the qualification above, carries direct and material consequences for portfolio construction.

For these reasons, VIX cannot be a reliable investable asset on the longer term for portfolio construction and risk diversification due to its expensive and unpredictable price dynamic. Which raises the question of uncertainty about risk itself as a distinct, independently priced dimension of financial risk, which is the conceptual foundation upon which this paper's central argument rests as VIX is not a trustworthy asset on the long-term to properly hedge or diversify.

Furthermore, VIX is not the only issue when it comes to trying to find the right asset to diversify risk, indeed Stambaugh, Yu, and Yuan (2015) [10] demonstrate that idiosyncratic volatility generates asymmetric arbitrage impediments across the mispricing spectrum, with short-sale constraints allowing overpriced, high-volatility stocks to remain chronically mispriced, a finding that establishes volatility not merely as a measure of return dispersion but as a structural determinant of market efficiency and return dynamics as these assets are considered too risky, and cannot be trusted.

Similarly, Hsu and Murray (2007) [11] provide an early empirical corroboration of this dynamic, showing that short-term changes in the VIX exhibit no significant correlation with subsequent realised volatility of the SPX, suggesting that the market's own uncertainty about future volatility is a distinct and imperfectly priced phenomenon from volatility itself which suggests Volatility and VoV are orthogonal, which we will confirm later on. Then, instead of using VIX, Baltussen, van Bekkum, and van der Grient (2018) [12] extend this logic decisively by constructing a market-based measure of risk ambiguity, the volatility of expected volatility extracted from option prices and demonstrating that stocks characterised by high uncertainty about their own risk profile underperform their low-ambiguity counterparts by 8% per year, driven not by conventional risk compensation but by a preference-based withdrawal of pessimistic investors that leaves affected assets systematically overpriced. In other terms as explained above, stocks with unpredictable volatility patterns (high VoV) would tend to underperform predictable volatility patterns (low VoV). And unfortunately, DR only uses stocks' volatility to compute diversified portfolios, not VoV. It is not able to assess stocks' VoV which ends up in blindly investing in high VoV or low VoV stocks without considering their impact on the portfolio's risk factor exposures which is, again, measured differently than the usual portfolio volatility formula as we will see later on.

Indeed this Volatility of Volatility measure is a complete different priced factor than normal volatility as Huang, Schlag, Shaliastovich, and Thimme (2018) [13] formalise this into a structural asset pricing argument, establishing that volatility of volatility constitutes a

significant independent risk factor carrying a negative market price, investors willingly accept negative expected returns on instruments that hedge against volatility uncertainty, confirming that the distribution of risk itself, and not merely its level, is a priced quantity. Which Branger, Hülsbusch, and Kraftschik (2018) [14] refine further by demonstrating that the slope of the VVIX term structure, capturing the differential pricing of volatility uncertainty across maturities, constitutes a distinct and superior predictor of option returns relative to the VIX level alone, with its informational content shifting systematically between continuous volatility-of-volatility and jump risk depending on prevailing market conditions.

Taken together, this establishes that risk is neither static nor one-dimensional: it is a layered, dynamic, and investor-relative phenomenon whose true structure is revealed only through factor-level decomposition, whose measurement is subject to its own uncertainty, and whose uncertainty is itself priced by financial markets, which is precisely the proposition this paper sets out to examine. And while DR measures how well return is diversified among a portfolio of assets, VDR measures how well risk is diversified across volatility sources (or risk factors in other terms) and it would complement the lack DR has measuring the actual source of risk by finding stable or unstable stocks which have uncorrelated volatility dynamics.

c. (Re)Defining Risk

Defining "risk" is critical to fully understand what this paper is about and how this concept is frequently mischaracterised in conventional investment practice. As seen above, risk can have different prisms, variance of returns, factor exposure, portfolio concentration…

In a portfolio, to summarise the very root of the definitions seen earlier, "risk" is the uncertainty of future portfolio returns, quantified by the entire distribution of potential outcomes, with a special focus on adverse outcomes (losses). Thus the primary measure used by every investor is volatility, calculated by the standard deviation of returns of the portfolio. Volatility gives a magnitude of uncertainty around which returns of the portfolio will be in the future.

Precisely, the source of this uncertainty resides in the stochastic nature of each constituent asset's return-generating process. Indeed, the portfolio's return is based on its constituents' return, and each of its assets have also their own volatilities, deriving from their own uncertainties about the return they are going to generate. Assets have volatility because their expected future value is constantly changing due to a continuous influx of new information, uncertainty, and shifts in investor behaviour, thus, in order to eliminate portfolio risk, an investor can either try to find an asset which will instantly reach the desired return aimed at, meaning the investor needs to buy an asset that returns him immediately the return wanted (an asset acquired at price P_0 and instantaneously liquidated at $P_1 > P_0$, yielding a deterministic return with zero variance, which by construction yields a return variance of zero). Or the investor can try to buy assets so that their weighted average return converges towards the aimed return *while* having their correlations, weights and volatilities perfectly matched so that portfolio volatility equals 0 (e.g. taking 2 assets with a correlation of -1, the same volatility and the same weight will give mathematically a portfolio volatility of 0).

Now, these two cases are extreme examples that do not appear often in real world. The first approximates arbitrage, a nominally risk-free operation, yet arbitrage is inherently transient: pricing discrepancies close as capital exploits them, returns are not always known in advance, and implementation costs frequently erode expected gains. The risks inherent to arbitrage strategies are therefore primarily operational in nature, encompassing execution costs, infrastructure constraints, and human error, rather than the asset-level return uncertainty that characterises conventional portfolio risk.

The second example approaches to what reality looks like, building a portfolio of assets to converge towards an expected return. However in this case, being able to find more than two assets with perfectly negative correlation (-1) is mathematically impossible. Which leaves the investor with no other choices than to accept volatility as an irreducible structural feature of any portfolio of risky assets. Assets will have their own volatilities and impact the global portfolio's volatility.

Inherently both of these methods embed risks to generate returns, they are being used as an illustration to expose two kinds of risk a portfolio is confronted to indeed two interpretations can be seen :

- For example, solely investing in an arbitrage strategy leaves the investor exposed to the risks to put in place the arbitrage, if the entity behind fails, the investor's portfolio falls with it. This can be seen as a "systematic risk" *for that investor's portfolio*, one risk factor on

which all a portfolio is exposed to can trigger the downfall of it. If the investor does not have the capacity or does not want to allocate to other diversifying strategies, the only way to protect against it is to maintain a contingency allocation or seek explicit tail-risk protection, such as option-based hedging strategies or reinsurance arrangements, in the event of counterparty insolvency.

- On the other hand, investing in a portfolio of a few assets would expose the investor to the asset's own individual risks (like bankruptcy) a way to avoid this risk, could be, if possible, to add more assets to the portfolio to diversify away the risk and avoiding being exposed to one or two single assets having a systematic risk for the portfolio. In this case the risk is no longer systematic but becomes diversifiable risk. Indeed diversifying by adding more assets reduces each asset's own systematic risk (e.g. bankruptcy) but does not necessarily imply that the portfolio is effectively diversified in a factor-risk sense as assets can have a common underlying risk factor. **In order to truly diversify, correlation of volatilities need to be considered**.

To sum up both risks in financial words, **"Systematic risk"**, *at a portfolio level,* is the volatility a portfolio is exposed to when all assets have the exposure to the same risk factor. Systematic risk is the moment when risk; **cannot be diversified away in a portfolio solely by adding uncorrelated assets**, it needs to be properly hedged away (e.g. Reinsurance/Insurance, Reducing Allocation, Option strategies,…). The baseline for Systematic risk is about acknowledging the risk and taking appropriate measures to protect against it.

On the other hand, **"Diversifiable risk" (Or Idiosyncratic risk)**, is the volatility a portfolio is exposed to when assets are not exposed to the same risk factor but to each asset's own return uncertainties. Diversifiable risk can be diversified away by adding assets which volatilities are uncorrelated to assets that the portfolio already has (e.g. adding Bonds to a portfolio of Stocks, adding consumer defensive to a technology portfolio,…) in order to stabilise a portfolio's volatility.

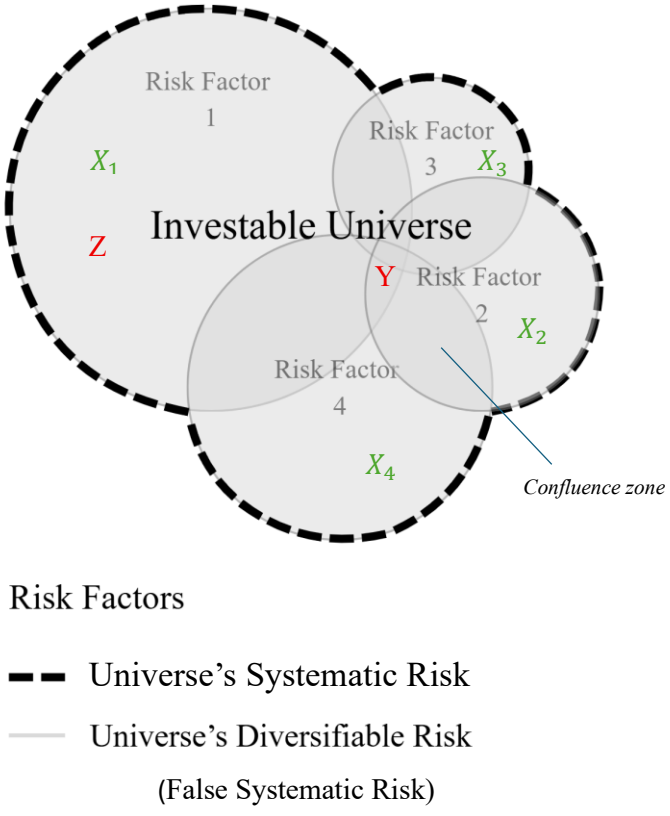
C) Conclusion

However a problem emerges, how to truly distinguish between a systematic risk and a diversifiable risk ? How to distinguish between a portfolio which is under-diversified and exposed a single diversifiable risk factor - and could be wrongly considered as systematic risk - given its investable universe and a portfolio truly hitting its diversifying limit being exposed to systematic risk ? **This is why it is critical to distinguish between True Systematic risk and False Systematic risk.** Where False Systematic risk stems from an under-diversified portfolio given its investable universe, solely exposed to a risk factor which could be diversified away from. False Systematic risk is in reality a Diversifiable Risk. True Systematic risk is what a portfolio is solely exposed to if it has sufficiently diversified away from its Diversifiable Risk. Leaving nothing more than the underlying Systematic risk factor of the investable universe itself (see figure c. below).

Indeed, Ilmanen and Kizer (2012) [15] demonstrate that conventional asset-class-allocated portfolios, despite their apparent breadth, mask a severe and unmitigated concentration in equity market directional risk, precisely the condition this paper identifies as False Systematic risk, whereby a portfolio appears diversified within its nominal structure while remaining entirely exposed to a single latent risk factor. Page and Taborsky (2011) [16] corroborate this finding through a regime-based framework, showing that asset class correlations spike dramatically during market contractions as indirect equity beta renders nominally distinct asset classes functionally identical in their factor exposure, a manifestation of the confluence zones identified in Figure a. below. **Roncalli and Weisang (2016) [17] formalise this argument most rigorously, demonstrating that a seemingly balanced asset-based risk parity portfolio can concentrate upwards of 80% of its real volatility into a single latent factor, and that only explicit factor-level risk budgeting restores genuine diversification**. Which is precisely what risk factor concentration figures below are representing and what issue this paper will try to tackle. Then, Hjalmarsson (2011) [18] extends this logic to the cross-section of equity characteristics, confirming that the classical diversification benefit, reduction of variance through imperfect correlation, applies identically to factor exposures, with diversification across orthogonal characteristics yielding substantially superior and more stable risk-adjusted returns than any single-characteristic allocation.

And as illustrated below (see figure a.), using Venn's diagrams, each region within the diagram represents the investable universe where, the surface is composed by assets which can be exposed to one or several risk factors. For instance, asset **Z** is exposed to risk factor 1 and asset **Y** is exposed to risk factor 1, 2 and 4. The false systematic risk would be having a portfolio of assets only exposed to a single risk factor – such as Z - while the investable universe allows for larger diversification. Similarly, a portfolio allocated to assets having a concentration of risk factors such as Y is also to avoid. As the concentration of risk factors increases, the assets' sensitivity to instability increases more easily than being exposed to several risk factors individually. Indeed, the diversified portfolio would be the portfolio which has assets exposed to all risk factors individually avoiding confluence zones in darker grey, an example would be to have a portfolio composed of (X_1, X_2, X_3, X_4). That portfolio would have individual exposures to risk factors which reduces the risk of assets' volatilities moving and spiking together, as one asset which experiences high volatility will not be followed by the other assets as they are exposed to different risk factors. Even though in reality assets will hardly only possess one risk factor but the main idea to retain is to truly diversify a portfolio one needs to make sure its portfolio is not solely invested in one risk factor or in a fully concentrated factor but be as much diversified as possible across them.

Figure a. Universe A's Risk Factor bubbles



Furthermore, it is critical to understand that the notion of risk is also highly relative depending on the investor's aim and constraints. For instance, a portfolio manager having a mandate to manage only US Tech Stocks will have a systematic risk factor that is "Tech Fear", however for a portfolio manager that has a larger mandate (e.g. US Stocks), it will not see Tech Fear risk factor as a systematic risk but as a diversifiable risk as it can diversify with other sectors. **At a portfolio level what is systematic risk for some can be diversifiable for others.** As illustrated below (see figure b.), taking two Investable universes, A and B. Universe A is taken from the previous illustration, composed by multiple risk factors (1, 2, 3, 4) outlined by the black and the dark red border. Universe B which is outlined by the red and dark red border, is bigger than A and comprises additional risk factors (5, 6, 7) on top of A's. In this case, risk factors 1, 2, 3 and 4 share together A's Systematic risk factor, meaning every asset in A's investable universe has a common risk factor which it cannot diversify away from. However, Universe B, being larger in terms of risk factor availability, allows portfolios exposed to assets in Universe A to diversify away from A's systematic risk, as Risk Factor 5, 6 and 7 have assets outside A's universe.

Figure b. Universe A & B's Risk Factor bubbles

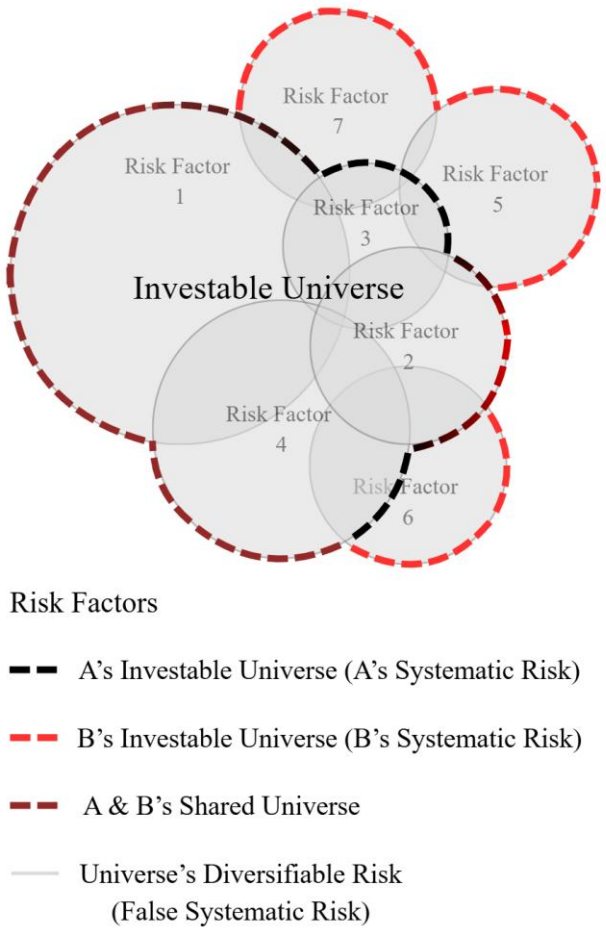


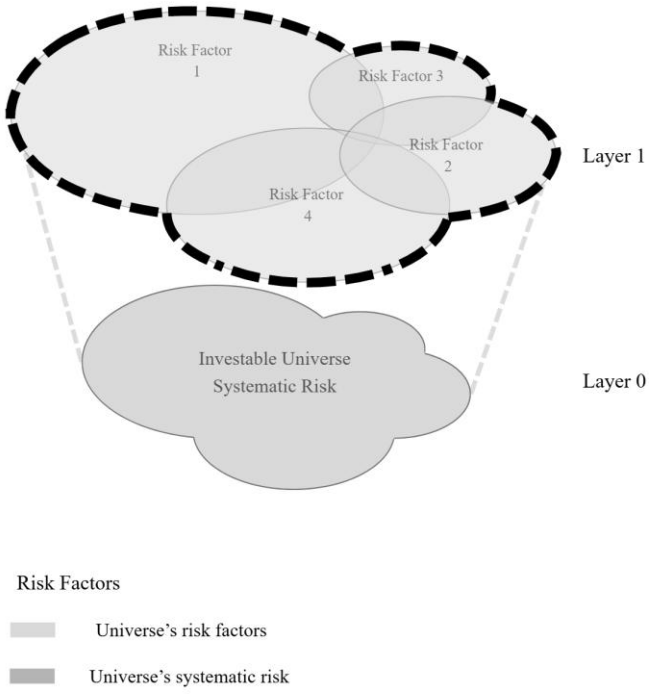
Figure c. models the risk factor layers, where risk factors from universe A have all in common an underlying risk factor which universe A itself. As a matter of example, let A be "US Stocks" and B "North American Stocks" (US + Canada). The different overlying risk factors could be identified as US Tech Stocks for risk factor 1, US Consumer Defensive Stocks for risk factor 2, US Energy Stocks for risk factor 3 and US Financial Services Stocks for risk factor 4. All these risk factors have an underlying risk factor which is being a US stock. If US Stock experiences a problem (US fear, geopolitical tensions involving US,…), all US stocks will experience market turmoil. On the other hand, being invested on Universe B (US + Canada) would offer better possibilities to reduce exposure to the US downturn and diminish the volatility from the simultaneous spike of the US stocks. As Canadian stocks here did not experienced the shock US stocks did.

Nevertheless, US and Canadian stocks also have an underlying common risk factor as they share the attribute of being North American Stocks, which is what can be called their "Layer 0".

Layer 0 should be considered as the very limit of a portfolio diversification possibility, as a portfolio cannot diversify itself away from its belonging universe. Layer 1 denotes the first stratum of decomposable risk factor exposures within the investable universe, across which portfolio diversification remains feasible, where a portfolio can diversify several risk factors. This hierarchical decomposition exhibits a fractal structure, replicable at successively finer levels of granularity. Indeed, within US Tech Stocks, there are "layer 2" risk factors that can be identified by narrowing the risk factor analysis (US Tech Large Stocks, US Tech Small Stocks,…). Additionally, even other risk factor clusters can be identified following other than an economic sector classification, (e.g. World-scale based : Continent-wide, Nation-wide,…).

A critical qualification pertains to the dynamic nature of factor sensitivities. As market structure changes (correlations between assets, trendy sectors,…) over time, asset's sensitivity to risk factors do as well. **Thus, what is represented by the bubble charts is the situation at time t. At time t+1 bubbles' size and interrelation may change as assets correlations change over time as well, meaning an asset sensitive to only one risk factor at time t, might be sensitive to two risk factors at time t+1.**

Figure c. Universe A's Risk Factor Layers



Furthermore, it is crucial to understand, risk factor does not measure volatility, it measures the sensitivity to a risk. Two assets can be exposed to the same risk factor but have different volatilities. And nothing impedes an asset to see its sensitivity to risk to change over time.

Thus to summarise, portfolio risk can be defined as the exposure to loss, whose magnitude, stability, and very nature as diversifiable or systematic are determined by the interaction between the assets' characteristics and the investor's constraints.

Now, the following sections tackle the method to achieve this true diversification, through the introduction of new metrics designed to measure volatility instability (VoV), volatility correlation (VDI) which will be used following the DR's methodology to spread risk factor risk in portfolio construction.

II) Introducing new ratios to measure Risk

A) Constructing Volatility Dispersion Index (VDI)

As stated in above, DR analyses the assets' return diversification based on their weighted average volatilities and based on the portfolio volatility. A possible idea to fix DR's lack of hindsight about risk diversification would be to consider correlation of assets volatilities' instead of the correlation of assets' returns. In other terms, instead of computing the portfolio volatility using the classic portfolio volatility formula (given below) which quantifies the portfolio's return volatility by using the variance-covariance matrix from the assets' returns, the idea would be to use the variance-covariance matrix derived from the change of volatilities (Volatility of Volatility). This would achieve a measure of the volatility of the portfolio's volatility, which would measure the stability the portfolio's risk over time.

Indeed, in classic Modern Portfolio Theory, portfolio volatility is calculated as such :

$$Portfolio\ Volatility = \sqrt{[w_i][\sum_{R_i}][w_i^T]}$$

Where $[w_i]$ *is the matrix of weights for i assets (*$[w_i^T]$ *is the transposed matrix)*
$[\sum_{R_i}]$ *is the Variance-Covariance matrix of assets returns*

In order to calculate the VoV of the portfolio, the steps are :

1) First, calculate the returns based on the assets' price change,
$$R_a = \frac{Price_{\ t}}{Price_{\ t-1}}$$
2) Then, the volatility of these returns either on a rolling window method or for all the period
$$\sigma = \sqrt{\frac{\sum_i^n (X_i - \bar{X})}{n}}$$
Or use any other formula on Python (stdev()) or Excel (=Stdev.s) to calculate directly the volatility
3) Compute the change period by period as if these were asset prices in step 1
$$\Delta_t = \frac{\sigma_t}{\sigma_{t-1}}$$
Where σ_t *is the volatility at time t*
Where Δ_t *is the change in volatility*
4) Afterwards, calculate the covariance matrix of these changes in volatility, here again, either on a rolling window or for all period
5) And finally, replace the covariance matrix of assets' returns in the portfolio volatility formula

This will give the **Volatility Dispersion Index (VDI)**, which is a measure of how synchronised volatilities of the portfolio's assets move.

$$Volatility\ Dispersion\ Index\ (VDI) = \sqrt{[w_i][\sum_{\sigma_i}][w_i^T]}$$

Where $[w_i]$ *is the matrix of weights for i assets (*$[w_i^T]$ *is the transposed matrix)*
$[\sum_{\sigma_i}]$ *is the Variance-Covariance matrix of assets' volatilities*

VDI calculates the volatility of volatility of the portfolio (stability), it tells if the assets' volatility tend to move together or not. The higher the index, the more likely the volatilities are to move together. It means the individual volatilities within the portfolio tend to rise and fall together, indicating a potential for unstable risk profiles during periods of market stress.

The lower the index, the less likely volatilities tend to move together suggesting the risk profile is stable across time, a low VDI suggests that the volatilities of the assets are diversified, meaning the overall portfolio's risk profile may be more stable even if asset returns are correlated.

The difference is that VDI studies a portfolio's stability by looking at risk sources, not return sources. Meaning that asset's returns can be correlated without having the same risk exposure. Stability is defined by the portfolio's ability to experience the same volatility over time. A portfolio experiencing a low but highly unpredictable volatility is more dangerous and insidious for an investor than a portfolio with a higher but more predictable volatility. Indeed an unstable portfolio gives a false sense of security to the investor who thinks the

low volatility portfolio will protect it against market turmoil.

VDI quantifies the Risk Model Risk as it tells if a 15% volatility-portfolio is reliable or not, it suggests if that 15% is going to be stable or unstable, meaning it could change drastically and go to 10% or 20%. Overall it's a volatility confidence indicator which measures how stable a portfolio is, suggesting how confident a manager can be with its risk management. To put it bluntly, when comparing two portfolios, a portfolio with a VDI of 5% is riskier than a portfolio with a VDI of 1% given it has more instability in its volatility.

VDI **does not replace** traditional risk metrics like Value-at-Risk or Portfolio Volatility, it rather suggests how unsure an investor can be about its risk management and how much risk stability it has in its current portfolio.

Even though, VDI is complementary to traditional risk management metrics (Portfolio Volatility, Drawdown, CVaR,…), it can also be used with returns by measuring return/stability efficiency. Indeed, by modifying the already-existing Sharpe Ratio by integrating the VDI of the portfolio instead of the portfolio volatility, the metric gives an alternative way to measure return/stability. This metric would measure alpha quality, or rather, the persistency and predictability with which alpha is produced by portfolio managers per unit of instability in the portfolio.

$$Stability\ Adjusted\ Alpha = \frac{R_p - R_f}{\sum_p}$$

Where
R_p *is the return of the portfolio*
R_f *is the risk-free rate*
$\sum_p$ *is the VDI (VoV) of the portfolio*

The Stability-Adjusted Alpha measures how persistent managers are able to generate returns across time with stable risk. A high Stability-Adjusted Alpha conveys a better efficiency in maximising return per unit of instability in the portfolio. A low SAA would mean the portfolio relies more on unstable risk to generate returns.

Maximising this ratio would build portfolios that generate returns from stable, predictable risk sources rather than unstable, systematic risk factor. Unlike Sharpe Ratio, which measures return per unit of total risk. It does not distinguish between good risk and bad risk, namely stable, predictable and diversifiable or unstable, unpredictable and systematic.

Also, SAA differs fundamentally from Treynor Ratio, as SAA does not measure market exposure (Beta) but measures how predictably and sustainably the alpha is being generated. Indeed Beta measures how an asset's return moves for a change in 1% of the market's return, with the same logic being applicable to a portfolio level. Thus Beta just tells whether the portfolio's return tend to move with the market or not, but it does not tell if those returns come from a stable or unstable source of risk. Where the Treynor ratio asks if the investor is getting paid for market exposure, measured by Beta, Stability-Adjusted Alpha asks whether the returns come from a diversifiable risk or a systematic risk.

This matters crucially as the three ratios measure fundamentally different elements, and have different use cases.

Ratio	Answers	Risk Perspective
Sharpe	Return per unit of Total risk	Quantity of Risk
Treynor	Return per unit of market risk	Market Risk
SAA	Return per unit of risk instability	Quality of Risk

The Sharpe Ratio measures excess return per unit of total risk, making no qualitative distinction between its components; it treats all volatility as equal, offering a purely quantitative assessment of a portfolio's risk.

The Treynor Ratio refines this by incorporating the Capital Asset Pricing Model (CAPM) perspective. It considers only systematic risk as the relevant metric, asserting that diversifiable risk should not be rewarded. Thus, it evaluates how efficiently a portfolio is compensated for its exposure to the market factor alone. However, this "market risk" assumes risk as a big and single systematic risk factor whereas there are several risk factors that can be outlined in reality. Market risk is determined by several underlying risk factors which – together – are considered as "market". Market can be seen a bundle of different risk factors, namely among others :

1. Economic Growth Risk (Will there be a recession?)
2. Inflation Risk (Will inflation spike or collapse?)
3. Interest Rate Risk (What will central banks do?)
4. Credit Risk (Will corporations default?)
5. Liquidity Risk (Can assets be sold easily?)
6. Volatility Risk (Will markets become unstable?)

Treynor Ratio measures how much excess return is generated by unit of market exposure, a bundled systematic risk **which is in reality diversifiable**. The notion of Beta is flawed for this reason.

In contrast, the Stability-Adjusted Alpha (SAA) introduces a fundamentally different dimension. It moves beyond the bundled market factor risk to assess its quality, meaning identifying the actual exposures to true systematic risk factors individually by studying how much correlated the volatilities are in the portfolio. By measuring return per unit of risk instability, it evaluates whether returns are generated from stable, predictable processes or unstable, erratic ones. This provides a qualitative judgment on the sustainability and resilience of the return stream itself.

B) Introducing Volatility Stability Ratio (VSR)

After calculating VDI, the next logical step is to integrate the index in the Diversification Ratio framework to calculate the baseline volatility stability of a portfolio. volatility stability indicates if the sources of risk of a portfolio are stable and predictable or unstable and unpredictable. As stated in introduction, Diversification Ratio gives the diversification of sources of return but not diversification of actual sources of risk. In practice, this means replacing the portfolio volatility by the VDI in the DR :

$$Volatility\ Stability\ Ratio\ (VSR) = \frac{\sum_{i=1}^{n} w_i \sigma_i}{\sum_p}$$

Where $\sum_p$ *is the VDI or the VoV of the portfolio.*

VSR measures how stable the source of volatility (risk) is for a portfolio given its average risk. The higher the ratio the more stable the volatility is from the sources of risk the portfolio is currently exposed to. Meaning for the given average level of risk taken, the portfolio's assets do not see their volatilities spike together, indicating that the portfolio's overall volatility exposure is stable and reliable.

The lower the ratio the more unstable the sources of risk are for the average level of volatility of the portfolio, indicating unpredictable risk profile that is prone to sudden spikes or drops given the portfolio's average risk. The portfolio's risk is highly unstable (high VDI). This means the volatilities of the assets move together, indicating a strong, common risk factor. Trying to reduce volatility level by adding more assets might be a waste of time. While this gives a false sense of low volatility in calm times, the portfolio's volatility will spike as all assets have a common risk factor when markets become volatile.

Compared to VDI, VSR allows to measure stability of risk compared to average risk in the portfolio. VDI is a standalone measure, it gives a part of the important information but does not compare it to a baseline. Sharpe Ratio and Stock return have the same relation, return is a measure on a standalone basis but it is not comparable on its own without considering volatility to achieve that return. For that reason, VSR is used as a baseline to compare portfolios VDIs. In that case, to get the highest stability of risk for a given portfolio, VSR has to be maximised, in other terms, VDI has to be minimised relative to weighted average volatility, which evaluates how much stability a portfolio can have given the innate riskiness of its assets.

VSR will punish "false" diversification, as the VSR will be low for portfolios that appear diversified in terms of returns but are concentrated in terms of risk factors suggesting unstable risk profile. A portfolio of 10 tech stocks might have a decent classic Diversification Ratio (DR) because the stocks aren't correlated, perhaps some tech companies are not exactly growth tech companies, like Apple or Microsoft (Both Large Core according to Morningstar classification) unlike stocks like Nvidia or Palantir (Both Large Growth still according to Morningstar). However, their volatilities will be highly correlated, which leads to a high VDI and therefore a low VSR. The VSR correctly identifies this as a poorly structured portfolio with an unstable risk profile.

VSR can also be considered as an indicator for market instability. As VSR decreases, volatility becomes increasingly more unstable, as correlation of underlying volatilities start increasing. Assets start moving together and the portfolio's risk factor concentration increases. This instability can be due to a bearish sentiment as well as a bullish sentiment. In both situations, whenever the market rallies or plummets, volatilities tend to spike.

C) Achieving Volatility Diversification Ratio (VDR)

Where VSR was answering the question of whether the portfolio had a stable risk profile or not for its average risk level, VDR answers the question of what kind of risk the portfolio is exposed to.

The last step in the framework is to integrate Weighted Average Volatility of Volatility (WVoV) in the VSR. The division of WVoV by VDI will give the Volatility Diversification Ratio (VDR). WVoV is calculated by

taking the volatility of the change in volatility of the returns. Basically taking the rolling volatility or the period volatility of data derived from step 3 above.

VDR measures volatility diversification and tells what is the source of instability of the portfolio, namely it tells if the portfolio's instability is driven by idiosyncratic risk or systematic risk :

$$Volatility\ Diversification\ Ratio\ (VDR) = \frac{\sum_{i=1}^{n} w_i \sum_i}{\sum_p}$$

Where $\sum_i$ is the volatility of volatility of the assets' portfolio

VDR measures risk factor concentration in a portfolio, it measures the degree to which the instability of a portfolio's risk is driven by idiosyncratic (diversifiable) events versus systematic (non-diversifiable) events, it directly tests for hidden exposure to a common volatility factor. A low VDR indicates that the assets' volatilities tend to spike together, signalling that a single, systematic source of uncertainty drives the entire portfolio's risk. Indeed, the VDR is not just sensitive to *if* volatilities co-move, but to the *impact* of that co-movement relative to individual instability. If the average individual instability (numerator) is high but volatilities are uncorrelated (leading to a moderate VDI), the VDR will be high, this indicates diversifiable instability. If the average individual instability is moderate but volatilities are highly correlated (leading to a high VDI), the VDR will be low which indicates non-diversifiable, systematic instability is the dominant component of the portfolio.

The key takeaway is that a high VDR means the source of the portfolio's risk is diversifiable, because, a priori in that situation, the current portfolio is already composed by assets which are not sharing a common risk factor, meaning the portfolio's risk can be decreased by adding more uncorrelated assets from the investor's current investment universe. An important point to note is that adding more diversifying assets will *increase* the VDR whatsoever as it means the portfolio's main contribution to risk is diversifiable. **VDR itself does not measure directly the quantity of risk in the portfolio, it tells which kind of risk is the primary driver of the portfolio's instability.**

The quantity of each risk can be approximated by :

$$Systematic\ Risk \approx VDI$$

Where Systematic risk can be interpretated as the Non-diversifiable risk which contributes to portfolio instability.
Total risk is given by the WVoV, as it represents the total risk if all volatilities had a correlation of 1, which the worst case scenario for the portfolio as diversification would not reduce portfolio's instability. Thus, Diversifiable risk can be approximated by the difference between Total risk and Systematic risk :

$$Diversifiable\ Risk \approx WVoV - VDI$$

In summary, Diversifiable Risk measures the actual amount of Systematic risk diversified away from the portfolio's instability. The higher the difference between the two, the bigger the risk amount has been diversified. If VDI = 5% and WVoV = 11% it means 6% out of 11% has been diversified away, or approximately 54% of the risk has been diversified.

To put it bluntly, VDR tells if it is worth trying to reduce the portfolio's instability or if it is a waste of time because the underlying risk is non-diversifiable**,** the problem is not the assets, it's their underlying risk factor (inflation, geopolitics, tech fear,...). If possible, the best way to reduce the risk is to add an asset which is not impacted by the common underlying risk factor, (e.g. For a portfolio of Stocks, adding Bonds).

When VDR is high, this suggests that there is a point in which the portfolio does not need to be diversified anymore which can be determined in several ways :

- VDR Plateaus: This is the primary signal. When adding a new asset no longer causes a meaningful increase in the VDR, the investor is no longer reducing systematic volatility risk
- VSR is High and Stable : The portfolio's overall risk profile has become stable and predictable, the "Stability Score" has been maximised
- Law of Diminishing returns : When the benefits of adding an additional diversifying asset dilutes returns, increases costs, adds management or monitoring complexities
- The target risk level has been reached

On the other hand, when there is a low VDR, portfolio risk is driven by a structural, non-diversifiable risk which cannot be hedged by adding more assets, but it does not mean risk is unescapable, the investor has solutions to hedge against systematic risk (e.g. Insurances/Reinsurances, Options strategies, Market neutral strategies, Futures,…).

To summarise, the table below synthesises all the indicators presented with their properties and the use cases.

	Primary Aim	Context of Use
VDI	Measures the instability of the portfolio's overall risk level.	Diagnostic tool to understand how much the portfolio's risk fluctuates over time.
VSR	Assess the stability and predictability of the portfolio's risk profile.	Used to determine if trying to lower volatility is useful
VDR	Diagnoses the nature/source of risk instability (diversifiable vs. systematic).	Used to determine if the portfolio's sources of volatility are systematic or diversifiable

When VDR is low, the portfolio is subject to systematic instability due to a concentration of the portfolio's assets on a few risk factors (e.g. Y in the confluence zones or Z only 1 risk factor from figure a.). In order to effectively fix the problem, incorporating uncorrelated volatility assets to break the correlation structure is needed. This will create a portfolio composed by assets individually exposed to independent risk factors (e.g. (X_1, X_2, X_3, X_4) in the figure c.).

VDR does not measure volatility per se, it measures if volatilities are stable and move together or not. The ratio tells when volatilities move together (high VDI) where portfolio is exposed to systematic risk, or on the other hand if volatilities tend to be independent from each other (low VDI). Meaning a high VDR could be the result of uncorrelated assets (low VDI) but with high volatilities. When VDR is already high, the objective is to stabilise the portfolio using the VSR. There, the aim is to lower volatility, by using VSR, the portfolio will maximise the stability efficiency given the average volatility of the assets by either allocating to lower volatility assets or directly allocating to stable assets with low VoV. Where VDR will consider the allocation on a global basis, by studying the interaction between assets, VSR will study the portfolio's assets individually to finetune their stability.

III) Building the Diversification Quality Framework

A) Diversification Dual Dimension Table

Just like VSR, VDR is orthogonal to DR, meaning these measures are complementary and combined they can provide a deep understanding of a portfolio's diversification exposures. The table below shows the dual diversification dimensions of a portfolio which is crucial to understand in order to be fully diversified in terms of returns and volatility :

Portfolio Diversification Dimensions	High VDR	Low VDR
High DR	The Robust Portfolio	The "All-Weather" Illusion
Low DR	The Predictable Bet	The Ticking Time Bomb

The Robust Portfolio : The portfolio is resilient, its diversification is likely to hold up during periods of market stress. Investor can have a high degree of confidence in its risk forecasts. This is the hallmark of a truly well-constructed, multi-asset class portfolio (e.g., a mix of global equities, government bonds, commodities, and alternative risk premia).

The "All-Weather Illusion" Portfolio : This is a portfolio that appears diversified but has a hidden vulnerability. It's the most insidious and dangerous quadrant. This diversification is a fair-weather phenomenon. In a true market crisis, the shared risk factor will dominate, correlations will converge to 1, volatilities will spike together, and the diversification benefit will vanish catastrophically (e.g : The portfolio of 4 different tech stocks above).

The Predictable Bet Portfolio : This is a concentrated portfolio, but with a stable and predictable risk profile. While the portfolio is exposed to the systematic risk of the chosen theme (portfolio loses money if the theme does poorly), investor can have high confidence in the *estimate* of that risk. The portfolio is unlikely to have surprise "volatility explosions" that are worse than expected. (e.g. A market-maker or a volatility arbitrage fund targeting a specific, stable amount of risk. They are concentrated in their strategy but have highly stable and predictable loss patterns.).

The Ticking Bomb Portfolio : This is a portfolio that is blatantly concentrated and obviously risky. This portfolio has no defence against a downturn in its specific sector or factor. It will experience extreme gains and extreme losses. Its risk profile is highly unstable and unpredictable. (e.g : A portfolio concentrated in a single sector, like energy or crypto using leverage. It is pure, undiluted, systematic risk.).

B) Proposing a Fully Diversified Portfolio

In order to build a truly diversified portfolio, the DR acts at the first level as the first diversifying barrier, looking if return sources are diversified, the VDR operates at second level measuring if the portfolio's instability is due to idiosyncratic risk or systematic risk,

VSR operates at the third level where it calculates if the portfolio's VoV is stable or not. And the final Diversification Quality Framework can be synthesized as the following:

Level 1 (DR): Are the returns diversified?
Level 2 (VDR): Is the instability diversifiable or systematic ?
Level 3 (VSR): Is the aggregate risk estimate stable?

However, a clear problem arises, with Level 2. In the case the portfolio's instability is due to a systematic risk factor, how should the investor know if this systematic risk due to a badly constructed portfolio which has been allocated to assets that - by coincidence or not - have all the same risk factor ? In other words, how to measure the "diversifiability" margin that the portfolio has given its universe ?

Indeed, when seeing the portfolio has a high systematic risk exposure after calculating VDR, a portfolio manager could be tempted to directly hedge - which may end-up increasing costs - without trying to first add other assets, as the current portfolio's systematic risk is a "false positive" and is in reality diversifiable as his universe still contains uncorrelated assets.

To manage this, the solution is to calculate the Diversifiability Margin Ratio (DMR) given by :

$$DMR\ (in\ \%) = \frac{VDR_p - VDR_{Benchmark}}{VDR_{UMax} - VDR_{Benchmark}}$$

Where VDR_p is the VDR of the portfolio, $VDR_{Benchmark}$ is the VDR of the benchmark used and VDR_{UMax} is the maximisation of the investor's universe VDR

The ratio is comprised between]-Inf;1], the portfolio not being able to be more diversified than the universe itself. This indicator gives two crucial information ;

- Whenever the DMR is negative, the portfolio concentrates more systematic risk than the benchmark, the portfolio is X% less efficient in diversifying risk.
- Whenever the DMR is positive, the portfolio is X% more efficient at diversifying risk than the benchmark

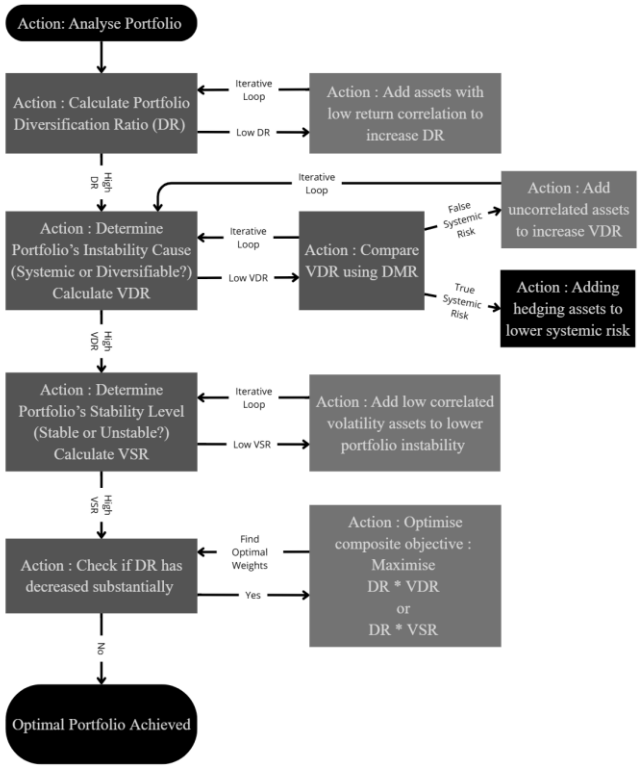
Even though, critical points need to be underlined, the universe's maximised VDR is a theoretical indication, as it could/should not be always attainable given its allocation might include illiquid stocks or could be concentrated in a few stocks (e.g. 4-5 stocks) which happen to have highly uncorrelated volatilities (which would expose to high diversifiable risk).

For these reasons, the right way of using the DMR is to compare with the portfolio's benchmark and then use marginal law of diminishing returns, using other indicators, or constraints the investor may have to balance benefits of adding diversification and increasing management costs (minimum number of stocks, liquidity,…). As a consequence, trying to reach 1 is not an optimal solution, but rather an indicator of how much margin there is to diversify in the current universe.

However in case the risk is truly systematic, the investor has to find a measure to contain potential risks from volatile markets. The portfolio then checks if it is either stable or unstable thanks to VSR before reassessing if DR has not decreased too much during the whole process. Once every step has been computed, the optimal diversified portfolio is reached.

The framework below, details the steps to follow to build a diversified portfolio. The framework remains flexible and can be easily integrated in a classic portfolio construction framework. Even the investor prefers skipping DR to go directly to VDR the framework allows it. Also, the investor is not compelled to maximise the ratios, as it can simply increase them compared to the initial portfolio which is sufficient to increase the stability and diversification of the portfolio.

Figure d. Proposed Fully Diversified Portfolio Framework



See Appendices for full size

In order to build a fully diversified portfolio, maximisation of DR and VDR has to be done. VSR is not being useful to actually understand the portfolio's sources of risk as it measures portfolio stability. However, the hard part of applying the framework above is that DR and VDR can have mutually exclusive goals. Several ways exist to maximise two distinct ratios. Depending on the priority given, either DR or VDR.

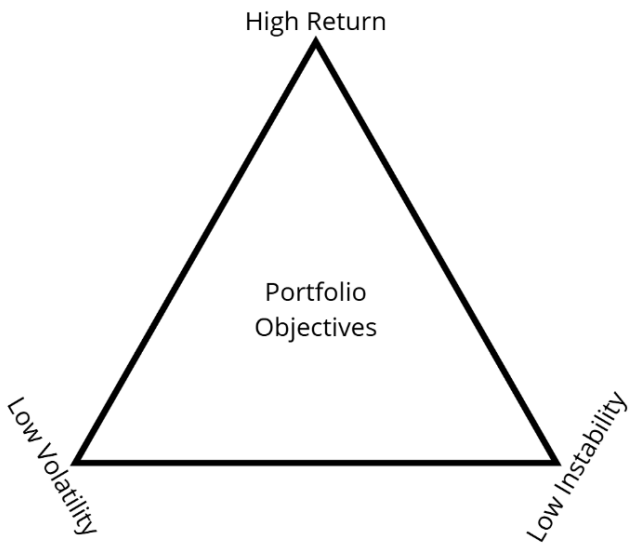
C) Nuancing the VDR

Realistically, neither VSR or VDR are miracle Ratio able to resist to every market shock, of course, their efficiency depends on the universe where it is applied. Presence of uncorrelated assets that have uncorrelated volatilities improves its efficiency. This framework helps improving risk management of portfolio managers when constructing their investments and choose if whether or not an asset should be considered or not.

An other blatant problem is the look-back nature of the indicator, VDR does not solve the problem of DR, being computed based on past data, thus not able to truly shape a portfolio to avoid market turmoil, even though VDR has provided insightful resilience results during tests, it is not sufficient to erase this drawback.

Additionally, the framework gives a good grasp of how a balanced and diversified portfolio should be built, and which portfolios should be avoided. The naïve idea would be to choose the Robust Portfolio just because it has a high DR and a high VDR. If the Robust Portfolio is supposedly stable, it can have either low volatility or high returns but not both. As reality is much more complex and this problem lies in the existence of the **Portfolio Trilemma** between High Return, Low Volatility and Low Instability.

Figure e. Portfolio Trilemma



A portfolio cannot have the three characteristics at the same time over long periods, investors need to choose among maximum two characteristics. Financial markets are generally efficient, if such a portfolio existed and was easy to identify, every investor would flock to it. This massive influx of capital would cause two things to happen :

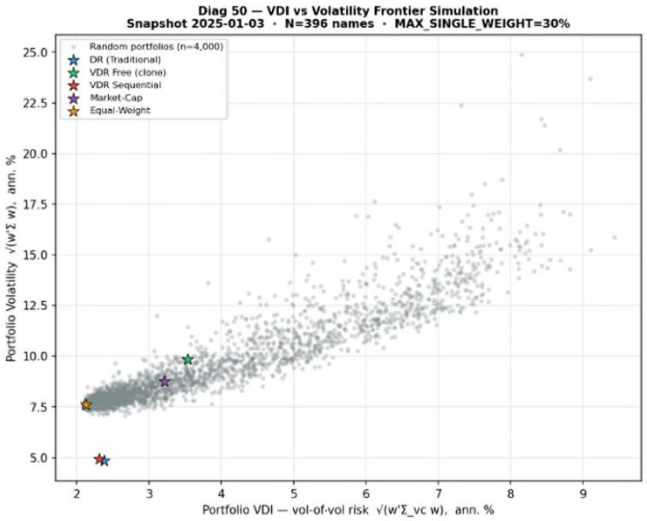
- Prices would be bid up, which lowers the future expected return, as investors have to pay more to enter the trade. The "High Return" characteristic would rapidly diminish.
- Correlations would increase, as all assets in the portfolio would start moving together driven by the same flows of capital. This could increase both volatility and instability.

Indeed, the perfect portfolio is a portfolio that has high returns with low volatility and low instability. By definition low instability implies low correlation with other assets' volatilities. And the pursuit of stability requires allocating to assets whose VoV movements are uncorrelated, often achieved through niche assets. These niche assets introduce a crucial trade-off which is the inclusion of assets with high individual VoV (e.g. Momentum or Emerging Market Equities) to benefit from uncorrelated instability, or assets with low VoV (e.g. Large Value or T-Bills) to serve as stable anchors. This allocation strategy is designed solely to maximize the stability of the portfolio's risk profile, but because VDI minimizes risk-of-risk rather than traditional risk magnitude, the resultant stability is often achieved by

sacrificing either low volatility or high potential returns. In other words, the "perfect portfolio" is inherently self-destructive. Its existence is a temporary arbitrage opportunity that the market quickly corrects.

Regarding VDR, by construction the ratio will push towards high VoV stocks with uncorrelated VoVs, this can only be achieved through more niche stocks, which can be small-caps or illiquid stocks. And this allocation inevitably increases volatility in the portfolio, and the relationship between Volatility and Stability plotted illustrate their efficient frontier. On the X-axis is plotted the portfolio VDI and Y-axis the portfolio volatility :

Figure f. Portfolio's Volatility-Stability Efficient Frontier



See Appendices for full size

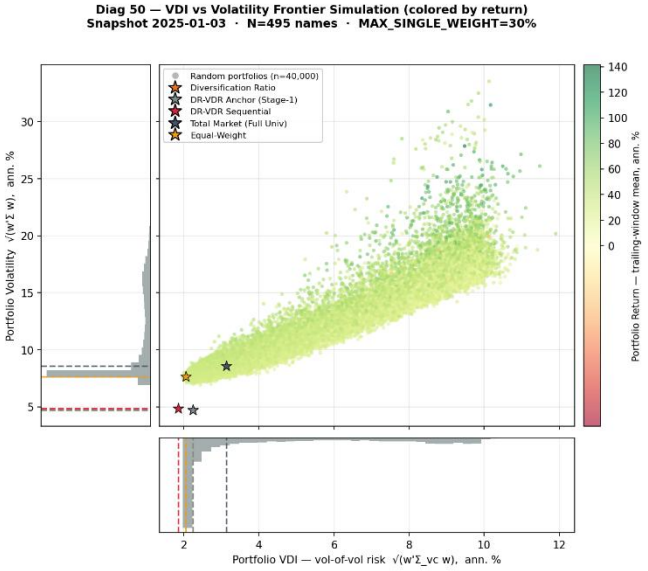
When plotting volatility against stability for 40,000 randomly simulated portfolios in a 400-stock universe on a graph, the efficient frontier shows a linear relationship, with efficient frontier being the line located beneath all the portfolios. As explained above, the portfolios cannot increase their VDI without accepting an increase in volatility as this means having to tilt towards more niche, small illiquid assets than usual.

However, there is not the existence of a "sweet spot" where the absolute peak of stability is achieved by combining different, uncorrelated *sources of risk* (like stocks and bonds), in this case, the investor directly has to choose whether to push towards more stability or less volatility.

In essence, whatever portfolio is above the frontier is not minimising its volatility for the stability given.

When adding return and distribution of portfolios for respective levels of volatility and stability :

Figure g. Portfolio's Volatility-Stability Efficient Frontier + Return

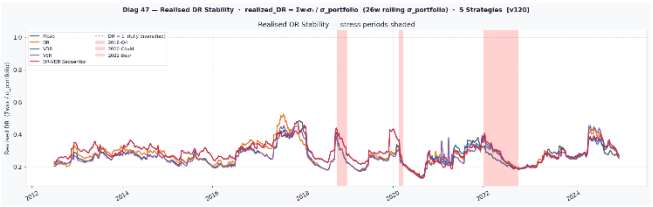


See Appendices for full size

For this graph, the same setup was used as for the previous graph *f*. The graph above shows a high increase from the bottom left to the top right. The upward-sloping curve of the frontier confirms the underlying financial dynamic: as a portfolio takes on more overall stability (X-axis), it becomes increasingly more difficult to maintain low volatility which is measured by volatility (Y-axis). The portfolios also look to have a double peak distribution where the first concentration is located at the bottom right and the second around the top right of the distribution with a reduced number of portfolios in the middle.

It is also critical to note that VDR and VSR do contain flaws in their investment approaches. As suggested earlier, increasing stability will tend to push towards niche assets that are uncorrelated from mainstream assets. Maximising VDR resumes to maximise Weighted VoV compared to VDI, which is tending to invest on niche assets that are likely to be highly uncorrelated assets that tend to exhibit instability in market turmoil, these can be smaller, less liquid assets for instance, which is all the contrary that should be done in order to avoid risk in a portfolio when market turmoil appears. This is why portfolio exclusively built around maximisation of VDR tend to underperform during downturns but outperform outstandingly in bullish markets. Even though, when using VDR in a reasonable window of optimisation, the model will help diversifying among all risk factors of the universe without pushing too far into niche risky assets.

Figure h. Portfolios' Realised DR



See Appendices for full size

As shown in figure h, the realised DR stability series, which measures how stable each portfolio's diversification ratio is over rolling windows, reveals both persistent differences across strategies and a clear structural shift over time. Contrary to what might be expected, DR and VDR move largely together over the full sample (correlation of 0.84), indicating that broad market-wide volatility regimes affect both strategies similarly, though each captures meaningfully distinct portfolio behaviour within those regimes.

Over the full 2012–2025 sample, VDR consistently runs at a higher realised DR level than DR (mean of 0.289 vs. 0.274), meaning VDR achieves a structurally higher diversification ratio on average despite not directly optimising for it. However, VDR is only below DR in 30% of weeks over the full sample, in the remaining 70% of weeks VDR sits above DR, which reflects the higher average level rather than any instability advantage.

There is a regime shift around 2018. Before that date, both strategies operated at higher average DR levels (DR at 0.294, VDR at 0.307), with VDR maintaining a consistent premium. From 2018 onward, both series compress downward (DR to 0.258, VDR to 0.275), though VDR retains its premium. Post-2022, the gap narrows significantly: DR averages 0.277 and VDR 0.273, and VDR is now below DR in 45% of weeks, close to parity, suggesting that VDR's structural diversification advantage over DR has largely eroded in the most recent period.

DR-VDR Sequential tracks VDR almost exactly throughout (correlation of 0.99, full-sample mean of 0.289), confirming that the DR floor constraint does not materially reduce the achieved diversification ratio relative to unconstrained VDR, while being below VDR in 54% of weeks, a slight but consistent drag attributable to the binding DR floor reallocating weight away from the highest-VDR positions.

And that is why the goal of the Diversification Quality Framework is not necessarily to find the mythical "perfect" portfolio, but to diagnose the trade-offs when considering an investment**.** The true value of analysing volatility instability via VSR/VDR is avoiding the seductive trap of a portfolio that *looks* perfect, namely with high return and low volatility but is actually taking hidden, unstable risks. It forces not only to question what the returns are, but also how sustainable and reliable the sources of these returns are.

Furthermore, even if the Robust Portfolio is considered as the natural choice for a manager, it should be considered more carefully. The Robust Portfolio maximises stability by allocating to both diversified sources of return and volatility which dilutes returns and ends up with a portfolio that is likely to underperform against Market-Cap weighted portfolios for instance.

All in all, out of the 4 portfolios proposed, the best choice to rely on is the Predictable Bet Portfolio. As it is the situation in which assets are selected according to an investing strategy or a vision from the investor. This investment manner is able to outperform the market if assets are wisely chosen. Indeed, this philosophy follows what is known today as "Stock Picking". Most of investors somewhat follow unconsciously this path as they invest in assets they think are going to outperform the market in each sector. For example, an investor aware of the basic rules of portfolio diversification knows it should not invest all in a single sector, thus it will buy other sectors (Tech, Energy,…). By doing so, it will naturally - but imperfectly - allocate to stocks which tend to not have the same risk factor, leading to a portfolio with low DR and somewhat a medium-to-high VDR. And in fact, the Predictable Bet Portfolio remains the best choice in terms of risk adjusted return, as the manager is aware of the correlation of its returns, and manages efficiently its risks by diversifying risk factors while avoiding too much return dilution, unlike the Robust Portfolio.

The idea behind VDR/VSR and VDI is to analyse and measure potential flaws that could be unknown to the investor. Risk is an integral part of investing and cannot be erased no matter how much a portfolio can be diversified. Markets do not fully behave in a rational and senseful way, no matter how mathematically two assets can be uncorrelated in theory. When investors start to fear the worse, the natural reflex is to sell everything by trying to anticipate what others will do in order to protect their own capital. In the end this mechanism increases all correlations.

Nevertheless, These indicators should be used to better assess risk management on a portfolio base. They are built to help investors understand how related their assets' risk factors are. Being able to measure risk allows better predictability and long term vision, the

idea behind the Diversification Quality Framework is to improve portfolio's resilience through appropriate decision making processes knowing whether the portfolio can/should be diversified or hedged instead.

IV) Backtested Portfolios

Before diving into the tests results and the implementation explanations, a small clarification is due.

There is indeed a wide gap between theory and practice, especially to transform a theoretical portfolio into a tradable strategy where adjustments are needed to solve real-world problems (illiquidity, size, delisting, IPOs, market regime shifts,…), the following part will detail the choices made in the backtest to enable the strategy to become tradable and to test the thesis of this research.

A) Dataset treatment & processing

a. Source and Coverage

The study uses the CRSP daily stock file accessed through Wharton Research Data Services (WRDS). The raw extract contains all common shares listed on NYSE, NYSE American (AMEX), and NASDAQ, identified by share codes 10 and 11, which restricts the sample to ordinary domestic common equity and excludes ADRs, REITs, closed-end funds, and other non-operating share classes. The full sample spans January 2010 to December 2025, but the empirical analysis begins in early 2012 to allow a two-year warm-up period for all rolling signal estimations (thus backtests effectively starting in 2012). Portfolio performance is therefore evaluated over the 2012–2025 window. The investable universe is fixed at the 500 largest stocks by market capitalisation at each weekly rebalancing date, analogous to an S&P 500 index profile in terms of liquidity tier and size mandate.

b. Delisting Adjustment

Delisting returns are handled in three steps to avoid the well-documented delisting bias in raw CRSP returns. First, wherever a numeric delisting return is available it is compounded directly into the last observed regular return for that security, so the terminal observation captures the full economic loss or gain at exit. Second, for securities that were involuntarily delisted for financial distress, identified by CRSP delisting status codes in the Shumway range, but whose delisting return is missing, a standardised imputed return of -30% is applied to the last trading observation, following the convention of Shumway (1997) [19] and Beaver, McNichols, and Price (2007) [20]. Third, ghost rows, daily observations that CRSP continues to carry after the true final trading date, are identified by the presence of any delisting record and are truncated, so stale post-delist data never enters return or volatility computations. When the delisting status code is absent from the extract, the Shumway imputation is suppressed and only numeric delisting return adjustments are applied; a warning is issued at run time to flag this limitation.

c. Weekly Frequency and Aggregation

All daily observations are aggregated to Friday-close weekly returns, yielding approximately 52 observations per stock per year. Within each calendar week, the compounded return is computed from daily log returns, requiring at least 2 valid trading days out of a possible 5 for a weekly observation to be retained. This threshold is permissive enough to retain holiday-shortened weeks while discarding stub observations where a stock was suspended or halted for most of the period. Market capitalisation and closing price are taken from the last available Friday observation of the week, with a fallback to the last available day of the week if Friday is missing. Returns are winsorised cross-sectionally at the 0.5th and 99.5th percentiles each week, meaning only the most extreme 1% of the return distribution, split evenly across both tails, is capped rather than removed.

d. Universe Construction

At each weekly rebalancing date, the eligible universe is formed through a sequential liquidity funnel. All stocks are first required to meet a minimum market capitalisation of $100 million and a minimum closing price of $5, both applied point-in-time to prevent look-ahead bias. Stocks passing both screens are then ranked by market capitalisation and trimmed to the largest 500 names. After the cap, any stock with fewer than one third of the required weekly observations within the current 60-week estimation window, that is, fewer than approximately 20 valid weekly returns, is dropped for that rebalancing date.

e. Volatility Estimation

Realised individual volatility is computed using an EWMA dispersion estimator with a span of 60 weeks, giving approximately 10 weeks half-life. This span was selected to balance responsiveness to regime changes against stability of the volatility estimate: a shorter span would react faster to volatility spikes but at

the cost of noisier week-to-week estimates, while a longer span would be smoother but slow to adapt during periods like the March 2020 sell-off. Weekly returns enter the estimator in levels and the output is annualised by multiplying by $\sqrt{52}$ to express volatility in annual units, consistent with the covariance matrix and all downstream signal construction.

The four active portfolios draw on two distinct covariance matrices, both estimated via the same EWMA Ledoit-Wolf procedure with a 60-week span. DR uses the return covariance Σ_{ret}, estimated on weekly log-returns. VSR and VDR use the vol-change covariance Σ_{vc}, estimated on weekly log-ratio volatility differences $\delta_{i,t}$. DR-VDR Seq requires both: Σ_{ret} enters the DR floor computation via its Cholesky factor L, and Σ_{vc} enters the Stage 2 VDR denominator constraint $y^{\top}\Sigma_{\text{vc}}y \leq 1$.

In all cases the estimation procedure follows an identical two-stage shrinkage pipeline. Stage 1 applies analytical Ledoit-Wolf shrinkage to the EWMA-weighted sample covariance, using a pseudo-data construction to accommodate the non-uniform observation weights. Stage 2 applies an additional adaptive shrinkage toward a scaled identity matrix with intensity:

$$\alpha = \text{clip}\left(0.5 + 0.10 \times \log\left(\frac{N}{250}\right),\ 0,\ 0.80\right)$$

At $N = 500$ this yields $\alpha \approx 0.57$. Without this two-stage procedure a 500×500 sample covariance matrix estimated from 60 weekly observations would be severely ill-conditioned, the number of free parameters exceeds the number of observations by a factor of more than 60, making direct inversion unreliable.

f. Matrix Alignment and Look-Ahead Protections

Several independent controls ensure that no information dated after the portfolio formation date enters the signal computation or covariance estimation at any stage of the pipeline.

The most fundamental protection governs how estimation windows are sliced. All rolling windows, for returns, volatility changes, and volatility-of-volatility, are cut to include only observations strictly before the current rebalance date. The return realised at the end of rebalance day is the return the portfolio is about to earn; including it in the covariance matrix used to form that portfolio's weights would constitute a direct look-ahead. An earlier version of the code used an inclusive upper bound on the window slice, introducing exactly this one-week contamination into covariance estimation, which has since been corrected.

Pre-computed signal matrices are extracted using a point-in-time lookup that retrieves the most recently available observation at or before the rebalance date t, never after. This matters because some signals are computed at irregular intervals, and a naive index-based lookup without the at-or-before constraint could silently retrieve a value from a future date if index alignment is not exact.

Finally, a minimum-history eligibility gate governs which stocks are permitted to receive signal-driven tilts. A stock's score is only allowed to influence portfolio weights once it has accumulated sufficient weekly return observations over a meaningful lookback period. This prevents newly entered or recently relisted stocks from receiving large tilts before enough history exists to estimate their signal reliably, without artificially excluding them from the broader universe construction earlier in the pipeline.

B) Signal Normalisation

The four active portfolios, DR, VSR, VDR, and DR-VDR Seq, all consume raw signal levels directly and are solved as true-ratio SOCPs. No z-scoring, soft-gate modulation, or quadratic utility reformulation is applied upstream of any optimiser. The purpose of this section is therefore to define the raw metric estimations that feed the optimisation engine: individual stock volatility $\sigma_i$and volatility-of-volatility ν_i.

a. Individual Volatility

Let $r_{i,t}$ denote the log-return of asset $i$during week t. Individual stock volatility is estimated row-by-row using an Exponentially Weighted Moving Average (EWMA) variance estimator with span $H = 60$ weeks, annualised by $\sqrt{52}$:

$$\hat{\sigma}^2_{i,t} = \sum_{\tau} \lambda^{\tau}\,(r_{i,t-\tau} - \bar{r}_{i,t})^2$$

where $\lambda = 1 - 2/(H+1)$ is the smoothing constant and $\bar{r}_{i,t}$is the rolling mean return. The annualised individual asset volatility is:

$$\sigma_{i,t} = \sqrt{\hat{\sigma}^2_{i,t}} \times \sqrt{52}$$

A minimum history requirement of $H_{\min} = 52$ non-missing rolling observations is enforced; assets breaching this threshold are assigned NaN and excluded from the optimisation universe for that week.

b. Volatility-of-Volatility

The volatility-of-volatility (ν) metric captures the structural dispersion of weekly log-ratio volatility changes:

$$\delta_{i,t} = \log\left(\frac{\hat{\sigma}_{i,t}}{\hat{\sigma}_{i,t-1}}\right)$$

Applying an identical $H = 60$-week EWMA window to these log-volatility differentials yields the raw VoV metric:

$$\nu_{i,t} = \sqrt{\sum_{\tau} \lambda^{\tau} \, (\delta_{i,t-\tau} - \bar{\delta}_{i,t})^2} \times \sqrt{52}$$

where $\bar{\delta}_{i,t}$is the rolling mean of the log-volatility differences. The $\sqrt{52}$factor ensures unit consistency with the annualised individual volatility entering the DR numerator.

c. Cross-Sectional Winsorisation

To insulate the cross-section against microstructural data errors, corporate action distortions, and extreme regime shocks, both $\sigma_{i,t}$and $\nu_{i,t}$are symmetrically clipped at the 5th and 95th cross-sectional quantiles before being passed to the optimiser:

$$\tilde{x}_{i,t} = \min\left(Q_{0.95,t}(x),\ \max\left(Q_{0.05,t}(x),\ x_{i,t}\right)\right)$$

where $Q_{p,t}(x)$denotes the contemporaneous p-th cross-sectional percentile. For the vol-change series $\delta_{i,t}$that feed the VoV estimator, a separate winsorisation pass at the same quantiles is applied before the EWMA dispersion step. The winsorised vectors $\tilde{\sigma}_t$and $\tilde{\nu}_t$are the direct inputs to all four optimisers.

C) Convex Optimisation and Portfolio Formulations

a. Convex Optimisation

The true fractional diversification ratio $\mathrm{DR}(\mathbf{w}) = (\mathbf{w}^T\boldsymbol{\sigma})/\sqrt{\mathbf{w}^T\boldsymbol{\Sigma}\mathbf{w}}$ is non-convex and numerically intractable under direct maximisation with long-only and weight-cap constraints. The optimisation engine exploits a standard change of variables to convert this into a convex Second-Order Cone Programme (SOCP), solved within a CVXPY environment using CLARABEL or OSQP.

Let $\mathbf{w}_t \in \mathbb{R}^N$represent portfolio weights at time t, with constraints $\sum_i w_{i,t} = 1$, $w_{i,t} \geq 0$, and $w_{i,t} \leq \omega_{\max} = 0.10$.

Note that ω_max = 0.10 here constrains the auxiliary variable y rather than the portfolio weight w directly; because the true per-stock cap is nonlinear under the y-transform, this is a deliberately loose in-solver proxy, and the actual binding weight limit (ω_max = 0.30 for DR/VSR/VDR, 0.05 for DR-VDR Seq) is enforced only after weight recovery, as described in Section IV.C.b

Market Capitalisation Portfolio (Mcap)

The benchmark is a plain float-weighted portfolio over the full ~500-stock pool, requiring no optimisation:

$$w_{i,t}^{\mathrm{Mcap}} = \frac{\mathrm{Cap}_{i,t}}{\sum_j \mathrm{Cap}_{j,t}}$$

No per-stock cap is applied; Mcap is the unconstrained float-wei0ghted reference.

SOCP reformulation

For DR, VSR, and VDR the fractional objective is linearised via the change of variables $\mathbf{y} = \mathbf{w}/\sqrt{\mathbf{w}^T\boldsymbol{\Sigma}\mathbf{w}}$, reducing to:

$$\max_{\mathbf{y}\geq 0} \boldsymbol{\phi}^T\mathbf{y} \text{ subject to } \mathbf{y}^T\boldsymbol{\Sigma}\mathbf{y} \leq 1, y_i \leq \omega_{\max} \cdot N$$

where $\boldsymbol{\phi}$ is the numerator signal vector and $\boldsymbol{\Sigma}$ is the denominator covariance matrix, both specific to each portfolio.

Weights are recovered as $\mathbf{w}_t = \mathbf{y}/\sum_j y_j$ and clipped to $\omega_{\max}$. All three portfolios operate on the full $pool_{500}$ no soft-gating or z-scoring is applied upstream of the SOCP. The numerator signals are computed directly from the raw return or vol-change windows passed into each optimiser call.

In all three cases, $\boldsymbol{\Sigma}$ is estimated via EWMA-weighted Ledoit-Wolf shrinkage (lw_covariance_ewma, span = 60weeks), annualised × 52.

Diversification Ratio Portfolio (DR)

$$\boldsymbol{\phi}^{\mathrm{DR}} = \boldsymbol{\sigma}_t^{\mathrm{ann}}, \boldsymbol{\Sigma}^{\mathrm{DR}} = \boldsymbol{\Sigma}_t^{\mathrm{ret}}$$

where $\sigma_{i,t}^{ann} = std(\boldsymbol{r}_i) \times \sqrt{52}$ is the plain (non-EWMA) standard deviation of the 60-week return window, annualised.

The denominator $\boldsymbol{\Sigma}_t^{\text{ret}}$ is the EWMA-LW covariance of weekly returns. The optimisation therefore maximises the ratio of aggregate individual return volatility to portfolio return volatility, the classical Choueifaty Most Diversified Portfolio.

Volatility Stability Ratio Portfolio (VSR)

$$\boldsymbol{\phi}^{\text{VSR}} = \boldsymbol{\sigma}_t^{\text{ann}}, \boldsymbol{\Sigma}^{\text{VSR}} = \boldsymbol{\Sigma}_t^{\Delta \ln \sigma}$$

The numerator is identical to DR annualised individual return volatility. The denominator is replaced by $\boldsymbol{\Sigma}_t^{\Delta \ln \sigma}$, the EWMA-LW covariance of weekly log-ratio vol changes, annualised $\times 52$. VSR therefore maximises individual return volatility exposure per unit of joint vol-change co-movement, rewarding stocks whose volatility level is high but whose vol dynamics are mutually independent.

Volatility Diversification Ratio Portfolio (VDR)

$$\boldsymbol{\phi}^{\text{VDR}} = \boldsymbol{\nu}_t^{\text{ann}}, \boldsymbol{\Sigma}^{\text{VDR}} = \boldsymbol{\Sigma}_t^{\Delta \ln \sigma}$$

The numerator is now the vector of annualised individual VoV levels $\nu_{i,t}^{\text{ann}} = \text{VoV}_{i,t} \times \sqrt{52}$, where $\text{VoV}_{i,t}$ is the last observation of the 60-week EWMA dispersion of $\Delta \ln \sigma_{i,t}$. The denominator is the same vol-change covariance as VSR. Annualising the numerator by $\sqrt{52}$ ensures unit consistency with the annualised denominator. VDR maximises aggregate idiosyncratic vol-of-vol exposure per unit of joint vol-change co-movement, isolating assets whose volatility paths are both active and mutually decoupled.

Sequential DR-Floor Volatility Diversification Ratio Portfolio (DR-VDR Seq)

DR-VDR Seq solves a self-contained SOCP each week that maximises the VDR objective subject to a floor on the portfolio's diversification ratio.

Stage 1 : DR floor computation. Rather than solving a fresh DR optimisation, Stage 1 evaluates the DR ratio achieved by the pre-computed DR-VDR-Free anchor weights w^{DR}, restricted and renormalised onto the intersection universe $\mathcal{U}_t$of stocks with valid return, VoV-covariance, and individual-volatility data:

$$\text{DR}_{\max} = \frac{\sigma^{\top} w^{\text{DR}}}{\sqrt{w^{\text{DR}\top} \Sigma_{\text{ret}}\, w^{\text{DR}}}}$$

This is a proxy for the achievable DR on $\mathcal{U}_t$, not a re-optimised bound. The DR floor is $\delta = \alpha \cdot \text{DR}_{\max}$with $\alpha = 0.85$. No weights, signals, or proximity constraints carry over from Stage 1 to Stage 2; the link is purely through this scalar.

Stage 2 : DR-floor VDR maximisation. The allocation is derived from scratch each week via the change of variables $y = w/(\mathbf{1}^{\top} w)$, exploiting the fact that both the VDR ratio and the DR floor are individually scale-invariant in y. The full programme is:

$$\begin{aligned} &\max_{y \geq 0} \nu^{\top} y \\ &\text{s.t.} y^{\top} \Sigma_{\text{vc}}\, y \leq 1 \\ &\sigma^{\top} y \geq \delta \cdot \| L^{\top} y \|_2 \\ &0 \leq y_i \leq \bar{\omega} \cdot n \forall i \end{aligned}$$

where $\nu$is the vector of annualised per-stock VoV ($\nu_i = \text{VoV}_{i,t} \times \sqrt{52}$), Σ_{vc} is the EWMA-LW covariance of the VoV series, σ is the vector of annualised individual return volatilities, $L = \text{chol}(\Sigma_{\text{ret}})$ with $\Sigma_{\text{ret}}$the EWMA-LW return covariance, and $\bar{\omega} = 0.05$ is the per-stock weight cap. Fixing the VDR denominator at unity converts the fractional objective into a linear maximisation over y, recovering the same optimum as maximising the nonlinear ratio directly over w. The DR-floor constraint is a genuine second-order cone inequality constraining the portfolio's diversification ratio, not its distance from any weight vector.

b. Portfolio Constraints and Weight Bounds

The constraint structures differ between the two portfolio families solved via optimisation: the SOCP portfolios (DR, VSR, VDR) and the DR-VDR Seq two-stage SOCP. The Mcap portfolio requires no optimisation and is described last.

DR, VSR, VDR — SOCP

These three portfolios are solved via a Second-Order Cone Programme formulated in the auxiliary variable $\mathbf{y} = \mathbf{w}/\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}$. The constraint set imposed during the solve is minimal:

$$\mathbf{y}^T \boldsymbol{\Sigma} \mathbf{y} \leq 1, y_i \geq 0, y_i \leq \omega_{\max} \cdot N$$

The box constraint $y_i \leq \omega_{\max} \cdot N$ is a loose convex approximation of the per-stock cap, since the

exact cap $y_i \le \omega_{\max} \sum_j y_j$is nonlinear in $\mathbf{y}$. Weights are recovered as $\mathbf{w}_t = \mathbf{y} / \sum_j y_j$, then hard-clipped to $\omega_{\max} = 0.30$ and renormalised. There is no relative band around market-cap weights, no active-share constraint, no turnover penalty, and no factor-neutrality constraint in any of these three optimisations.

DR-VDR Seq — two-stage SOCP

The Stage 2 constraint set imposed during the solve is:

$$y^\top \Sigma_{\text{vc}}\, y \le 1, \sigma^\top y \ge \delta \cdot \| L^\top y \|_2, 0 \le y_i \le \bar{\omega} \cdot n$$

There is no simplex constraint inside the solver. Portfolio weights are recovered post-solve as $w = y / \mathbf{1}^\top y$, then hard-clipped to $\bar{\omega} = 0.05$and renormalised. Turnover control is delegated entirely to a 5% post-solve buffer band; there is no in-solver turnover penalty, no Euclidean ball around an anchor, no active-share constraint, and no factor-neutrality constraint. The solver chain is CLARABEL with SCS as fallback; OSQP is excluded as it cannot represent second-order cone constraints.

c. Buffer-Band Rebalancing Width and Market-Cap Fallback Anchoring

Two further mechanisms operate outside the optimisation objectives themselves, applied identically as post-solve steps across all five portfolios, and are responsible for the bulk of week-to-week turnover suppression and numerical robustness in the framework.

Rebalancing width (buffer band). After each weekly solve, the optimal weight vector $w_t^{opt}$is passed through a no-trade buffer before being adopted. For each stock p, the buffered weight is:

$$w_{p,t}^{buf} = \begin{cases} w_{p,t-1}^{buf} & \text{if } | w_{p,t}^{opt} - w_{p,t-1}^{buf} | \le \beta \\ w_{p,t}^{opt} & \text{otherwise} \end{cases}$$

followed by renormalisation,

$$w_t^{buf} \leftarrow w_t^{buf} / \sum_p w_{p,t}^{buf}.$$

The threshold β, the rebalancing width, is the absolute distance a stock's optimal weight must move from its previously *held* weight before a trade is executed at all; positions inside the band are left untouched rather than re-traded toward a marginally different target. This is the dominant turnover control for DR, VSR, VDR and DR-VDR Seq which carry no turnover term in their own SOCP objectives: a single global buffer $\beta = 5\%$ is applied to all three. As a guard against the band silently undoing the per-stock cap κ— a new entrant solved at the cap, $w^{opt} = \kappa$, $w_{t-1}^{buf} = 0$, would otherwise satisfy $| \kappa - 0 | \le \beta$ whenever $\beta \ge \kappa$and be frozen at zero, buffering is skipped entirely whenever $\beta \ge \kappa$.

Market-cap fallback anchoring. Independently of the buffer band, every optimiser in the framework shares a single default output for weeks in which no usable solution can be produced, the first rebalance date (no prior weight to carry forward or tilt around), a solver exception or infeasible/unbounded status, or a degenerate cross-section (a flat or all-zero signal). In all such cases the portfolio's weight vector is set to the contemporaneous float-weighted market-capitalisation portfolio over the relevant sub-universe:

$$w_{p,t}^{mcap} = \frac{mcap_{p,t}}{\sum_q mca\, p_{q,t}}$$

with equal-weighting, $w_{p,t} = 1/n$, substituted only in the edge case of zero aggregate market cap. This is not a Bayesian prior blended into any objective, DR, VSR, VDR and DR-VDR Seq all draw on the identical fallback routine, and it is the literal weight vector adopted whenever the optimiser itself fails to deliver one. The effect is that no portfolio in the framework can return a NaN or degenerate weight vector in any week regardless of solver behaviour: market-cap weighting functions as the structural floor that every strategy reverts to, rather than to zero exposure, a stale carry-forward, or an unrenormalised partial solution.

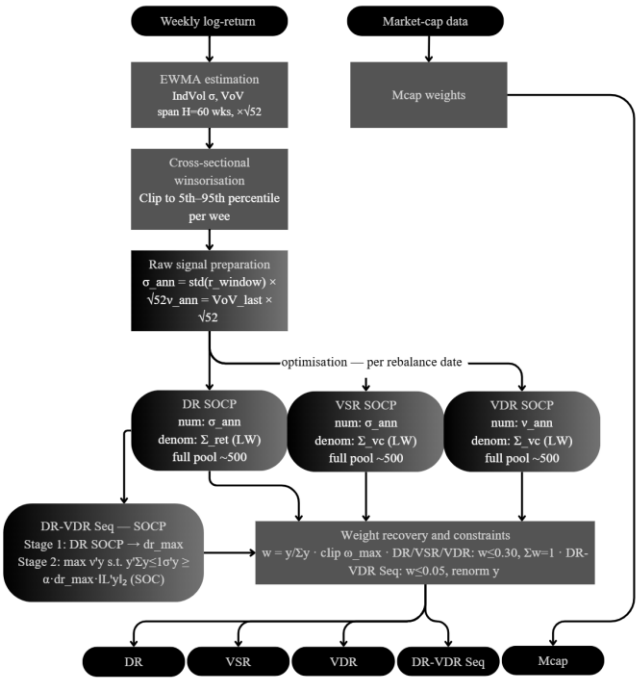
Marketcap-weighted

The market-capitalisation portfolio is constructed directly as:

$$w_{i,t}^{\text{Mcap}} = \frac{\text{Cap}_{i,t}}{\sum_j \text{Cap}_{j,t}}$$

over the full ~500-stock pool. No optimisation, no weight cap, and no constraint of any kind is applied. It serves as the unconstrained float-weighted reference portfolio.

Figure i. Portfolios steps



See Appendices for full size

D) Empirical Backtest Results and Performance Analysis

a. Consolidated Performance Overview

The following section presents a comprehensive evaluation of the five portfolio strategies over the 665-week backtest period spanning April 2012 to January 2025, covering full market cycles including the 2018 volatility episode, the 2020 Covid-19 market dislocation, the 2022 bear market, and the subsequent megacap-led bull market of 2023–2024. All performance figures are reported net of a 10 basis point one-way transaction cost assumption, applied to weekly rebalancing, the universe is restricted at each rebalancing date to the top 500 US equities by market capitalisation satisfying minimum price and liquidity filters stated earlier:

Figure j. Overall Results

Metric	Mcap	DR	VDR	VSR	DR-VDR Seq
Ann. Return (%)	**+14.67**	+13.72	+13.69	+13.18	+13.77
Ann. Volatility (%)	+15.90	+15.29	+14.39	+15.52	**+14.31**
Sharpe Ratio	0.923	0.897	0.951	0.849	**0.962**
Max Drawdown (%)	−31.10	−31.38	−25.23	−33.56	**−25.22**
CVaR 95% (bps)	−263.98	−255.58	−235.08	−262.20	**−231.85**
Ann. Turnover (%)	+70.25	+75.96	+137.22	**+21.15**	+131.28
Calmar Ratio	0.472	0.437	0.543	0.393	**0.546**
FF6 Alpha (%)	+0.08	**+0.89**	+0.58	−0.10	+0.65
FF6 t-stat	+0.52	+0.92	+0.39	−0.12	**+0.44**
Stab-α	138.9	126.56	157.23	119.69	**158.83**

Stab-α represents the Stability-adjusted alpha ratio exposed earlier. Consolidated performance metrics, April 2012–January 2025. All returns net of 10 bps one-way transaction cost. CVaR reported at the 95th percentile. FF6 alpha estimated via full-period OLS regression on Fama-French six-factor model. Highlighted cells (blue) indicate the best value in each metric row across optimised strategies.

Figure j. reports headline risk-adjusted performance statistics for the full backtest period. DR-VDR Seq achieves the highest Sharpe ratio (0.962) and the highest Calmar ratio (0.546) of any strategy in the group, including the benchmark, while recording the lowest CVaR at the 95th percentile (−231.85 bps). Together with VDR (Sharpe 0.951, Calmar 0.543, max drawdown −25.23%), these two strategies deliver

materially superior risk-adjusted outcomes relative to both the Mcap benchmark (Sharpe 0.923) and the DR optimiser (Sharpe 0.897). The co-optimisation architecture of DR-VDR Seq, anchoring the portfolio to a high-DR solution before applying the VDR tilt—is validated by its simultaneous leadership on Sharpe, Calmar, and CVaR.

The Mcap benchmark generates the highest gross annualised return of the group (14.67%), primarily attributable to its unrestricted concentration in the megacap technology sector during 2023–2024. All four optimised strategies sacrifice a portion of gross return, between 0.90 and 1.49 percentage points, in exchange for materially lower tail risk. The return shortfall is entirely accounted for by the post-2022 regime, as demonstrated in the sub-period analysis below; prior to 2023, every optimised strategy matched or exceeded Mcap returns on a risk-adjusted basis.

The Stab-α statistic, measuring return consistency through time, it ranks DR-VDR Seq highest (158.83), followed by VDR (157.23) and Mcap (138.93). DR scores 126.56, and VSR is last at 119.69, confirming that VSR's sub-par returns are not confined to isolated periods but represent a structurally inferior return profile across the full backtest horizon.

VSR warrants particular attention as a negative result. Its Sharpe of 0.849 is the lowest of any optimised strategy, its maximum drawdown (−33.56%) is worse than the benchmark, and its annualised FF6 alpha is negative (−0.10%, t = −0.12). However, its annualised turnover of only 21.15%, lower even than the passive Mcap benchmark, is anomalous and diagnostically important: it is not a sign of portfolio discipline but rather a symptom of the VSR optimiser converging to a near-equal-weight static solution, as confirmed by the weight and active share diagnostics presented

b. Sub-period and Regime-Conditional Performance

Even though in terms of overall performance DR-VDR-Seq tends to be the best risk-adjusted strategy, the sub-period analysis brings a more nuanced analysis.

Figure k. Portfolio NAV evolution



See Appendices for full size

A 26-week rolling Sharpe ratio was computed weekly for each strategy. DR-VDR Seq leads Mcap in 52.3% of rolling windows, the only optimised strategy with a majority win rate. VDR is close behind at 50.6%, while DR and Mcap are near-even at 49.2%. VSR leads Mcap in only 39.1% of windows, confirming its structural underperformance relative to the benchmark on a rolling risk-adjusted basis.

Figure l. decomposes performance into four market regimes. Several findings are noteworthy. First, in the 2012–2016 post-crisis bull market, both VDR and DR-VDR Seq achieve a Sharpe of 1.470 versus Mcap's 1.121, establishing the VoV signal's efficacy in an environment of normalising cross-sectional volatility dispersion. Second, in the 2017–2019 low-volatility expansion, DR-VDR Seq (1.426) and VDR (1.392) maintain leadership, while standalone DR (0.928) falls below the benchmark (1.228), suggesting that the DR objective alone is insufficient during periods of persistently low volatility dispersion. Third, and most critically, in the 2020–2022 stress and recovery cycle, DR achieves the highest risk-adjusted performance among optimised strategies (Sharpe 0.512), outperforming VDR (0.455), VSR (0.412), and DR-VDR Seq (0.433), and substantially outperforming Mcap (0.317). The defensive value of the DR anchor is most visible precisely in the period of maximum market stress.

The 2023–2025 megacap bull market presents a structurally adverse environment for all diversification-based strategies. Mcap achieves a Sharpe of 2.036, driven by concentration in a small number of large technology names that appreciated sharply during the artificial intelligence investment cycle. DR performs least badly among optimised strategies (1.548), benefiting from its broader weight distribution and partial overlap with the megacap universe through its near-passive active share. VDR (1.144) and DR-VDR Seq (1.265) trail more substantially, reflecting their concentrated active books and CMA factor headwinds documented in figure l. VSR (1.520) performs better than expected in this regime because its near-equal-

weight solution provides passive-like exposure to the index's market-cap concentration through index replication rather than by design.

Figure l. Sub-Period Sharpe Analysis

Sub-period	Mcap	DR	VDR	VSR	DR-VDR Seq
2012–2016 (Post-crisis bull)	1.121	1.299	**1.470**	1.123	**1.470**
2017–2019 (Low-vol expansion)	1.228	0.928	1.392	1.045	**1.426**
2020–2022 (Stress / Recovery)	0.317	**0.512**	0.455	0.412	0.433
2023–2025 (Megacap bull)	**2.036**	1.548	1.144	1.520	1.265

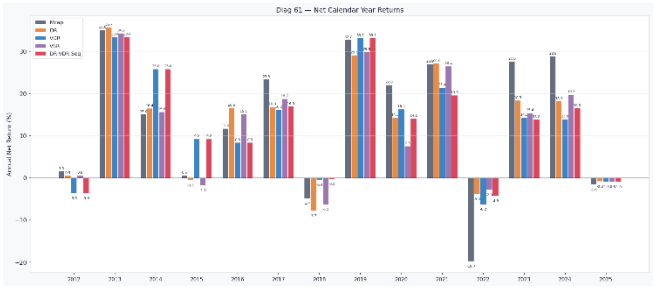


See Appendices for full size

Sub-period annualised Sharpe ratios. Best value per row highlighted in blue. Sub-periods are selected to capture distinct market regimes: post-crisis bull market, low-volatility expansion, stress/recovery cycle, and megacap concentration phase.

On a calendar year basis, DR, VDR, and DR-VDR Seq each win 6 of 14 years against Mcap; VSR wins only 4. All strategies record their worst relative year in 2024 (excess returns: DR −10.7%, VDR −14.9%, DR-VDR Seq −13.8%) and their best relative year in 2022 (DR +15.8%, VDR +13.5%, DR-VDR Seq +15.4%, VSR +16.9%), confirming a consistent pattern of outperformance during risk-off episodes and underperformance during megacap-led expansions. Mean calendar-year excess returns are negative for all strategies: DR −1.34%, VDR −1.30%, DR-VDR Seq −1.24%, VSR −1.97%.

Figure m. Calendar Year Returns



See Appendices for full size

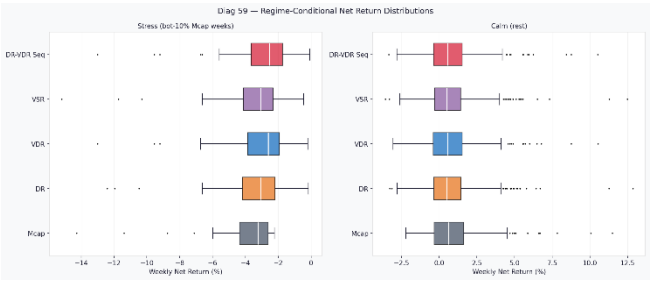
c. Stress-Regime Return Distributions, Return Asymmetry & Drawdown Analysis

Regime-conditional return distributions were computed by classifying each week as either a stress week (bottom decile of Mcap weekly returns, n = 67 weeks) or a calm week (n = 598 weeks). During stress weeks, DR achieves the lowest mean loss at −2.82% per week (median −2.35%), with a P5 return of −6.71%, confirming that the diversification-ratio anchor provides the most reliable capital protection in acute market dislocations. VDR and DR-VDR Seq record comparable stress-period means of −3.11% and −3.14% respectively, but diverge sharply on tail behaviour: DR-VDR Seq exhibits the least negative skewness of any strategy (−1.89), against VDR's −2.38% and Mcap's −2.93%, indicating that the DR anchor stage specifically suppresses extreme left-tail outcomes even when mean losses are similar. DR-VDR Seq's P5 of −6.14% and P95 of −0.25% confirm this: it records the shallowest worst-percentile loss and, uniquely, a near-zero 95th-percentile stress return, meaning that in roughly the best fifth of its stress weeks it essentially breaks even. VSR records the worst stress profile on every dimension: mean −3.80% (matching the benchmark's −3.88%), median −3.16%, P5 of −7.81%, and a skewness of −2.10, confirming that its near-passive convergence provides no stress-period differentiation from the market-cap benchmark.

During calm weeks, performance convergence is near-complete across optimised strategies. Mean weekly returns range from 0.595% (DR) to 0.683% (VSR), against Mcap's 0.749%. DR-VDR Seq (0.646%) and VDR (0.620%) sit centrally. The calm-period positive skewness of DR-VDR Seq is the lowest of the group (1.08 versus Mcap's 1.31), reflecting a smoother, less momentum-driven participation profile, consistent with a strategy that neither chases the largest up-weeks nor suffers the deepest down-weeks. The stress-period ordering DR > VDR ≈ DR-VDR Seq > VSR ≈ Mcap on

mean loss, and DR-VDR Seq > DR > VDR > Mcap on tail skewness, jointly constitute the primary empirical justification for the two-stage architecture the DR floor constraint delivers the best tail-shape improvement by preventing the VDR objective from concentrating risk into a small number of high-signal names during dislocations, while the VDR objective within the SOC-constrained feasible region maintains idiosyncratic diversification that pure DR cannot achieve, and the combination in DR-VDR Seq captures both properties simultaneously.

Figure n. Conditional Regime Distribution



See Appendices for full size

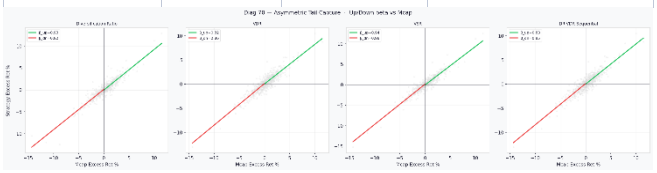
Figure o. reports upside and downside market betas, estimated from the subsample of 397 up-market weeks and 268 down-market weeks respectively (classified by the sign of the weekly Mcap return). The asymmetry ratio, defined as the downside beta divided by the upside beta, measures whether a strategy captures more upside than downside market exposure. A ratio below 1.0 is favourable.

DR is the only optimised strategy achieving a favourable asymmetry ratio (0.992): upside beta 0.928, downside beta 0.921, implying it loses marginally less in down markets than it gains in up markets relative to the benchmark. This result is consistent with DR's near-passive character, it tracks the market closely and does not systematically buy high-beta names in either direction. VDR's asymmetry ratio of 1.050 is the worst of the optimised strategies: it participates less in up markets (beta 0.819) than in down markets (beta 0.860), reflecting the VoV signal's tendency to maintain defensive positioning into recovery phases. DR-VDR Seq improves on VDR (ratio 1.026), with the DR anchor stage providing marginal upside participation, but the ratio remains above 1.0. VSR (ratio 1.048) suffers similar asymmetry despite its lower overall market beta, driven by its high downside beta (0.981) during stress weeks, consistent with its worst-drawdown-in-group finding.

The practical implication is direct: the Sharpe advantage of VDR and DR-VDR Seq over Mcap is achieved *entirely* through volatility reduction and drawdown compression, not through favourable upside participation. The strategies generate value by limiting losses more than they limit gains in absolute terms, but they do not capture disproportionate upside. This is an honest characterisation of the VoV approach's return source, and stands as the primary quantitative limitation to document when presenting these results.

Figure o. Return Asymmetry

Strategy	β Up-mkt	β Down-mkt	Asymmetry Ratio
DR	0.928	0.921	**0.992 (favourable)**
VDR	0.819	0.860	1.050 (unfavourable)
VSR	0.936	0.981	1.048 (unfavourable)
DR-VDR Seq	0.831	0.852	1.026 (marginal)



See Appendices for full size

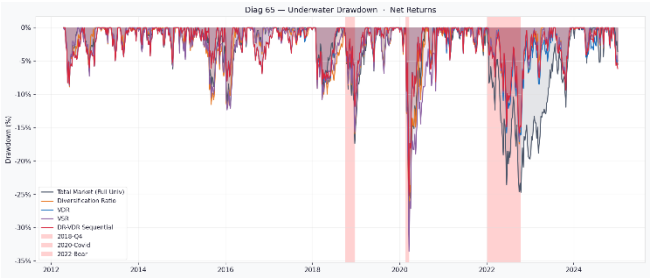
Figure o. Tail capture analysis: upside beta (β Up) estimated from 397 up-market weeks; downside beta (β Down) from 268 down-market weeks. Asymmetry Ratio = β Down / β Up. A ratio below 1.0 implies proportionally better protection than participation sacrifice.

All strategies record their maximum drawdown trough on 20 March 2020, coinciding with the acute phase of the Covid-19 market dislocation. The cross-sectional ordering of drawdown depth is informative: VDR and DR-VDR Seq achieve the shallowest troughs at −25.23% and −25.22% respectively, compared to −31.10% for Mcap, −31.38% for DR, and −33.56% for VSR. The 6-percentage-point drawdown advantage of the VoV-based optimisers over the benchmark is economically material for institutional mandates subject to drawdown-based risk constraints.

Recovery dynamics further differentiate the strategies. VDR and DR-VDR Seq both recover their pre-drawdown peak in 16 weeks, the joint fastest of the group. DR and Mcap recover in 20 weeks each, while VSR requires 34 weeks, the slowest. The joint drawdown depth and recovery speed of VDR and DR-VDR Seq is the most favourable risk characteristic observed across all evaluated dimensions.

The number of weeks spent below a 10% drawdown threshold quantifies sustained underwater exposure: DR-VDR Seq records the fewest such weeks (17), followed by VDR (23), DR (36), VSR (40), and Mcap (83). The large gap between the VoV-based optimisers and the benchmark on this metric confirms that their advantage is not limited to the peak drawdown date but extends through the recovery period, reflecting a structurally more stable NAV path.

Figure p. Drawdown analysis



See Appendices for full size

d. Portfolio Risk & Return Decomposition & Fama-French 6 Factors

Figure q. reveals the most structurally diagnostic finding of the analysis. DR and VSR exhibit mean active shares of only 32.4% and 30.4% respectively, implying that 68–70% of their portfolio weight is effectively passive relative to the market-capitalisation benchmark. Their mean tracking errors (4.89% and 4.06% annualised) are similarly modest. The standard deviations of both TE series are high relative to their means (DR: σ = 2.54%; VSR: σ = 2.59%), indicating that these strategies are not consistently active but instead oscillate between near-passive and moderately active configurations, hitting constraint boundaries intermittently rather than systematically pursuing their stated objectives.

By contrast, VDR and DR-VDR Seq maintain mean active shares of 77.0% with low time-series variability ($\sigma \approx$ 9.1–9.2%), reflecting genuinely and consistently active positioning relative to the benchmark. Their tracking errors of approximately 6.1% with standard deviations of only 1.37–1.38% indicate stable active risk taking. The near-identical active share and TE profiles of VDR and DR-VDR Seq, and their 0.993 weekly return correlation, confirm that the two strategies run the same active book. The marginal contribution of the DR-VDR Seq architecture relative to standalone VDR operates in risk-budget space (Figure q.) rather than return or weight space.

Figure q. Active Share & Tracking Error

Strategy	Mean AS %	Std AS %	Mean TE %	Std TE %
DR	32.4	11.4	4.89	2.54
VDR	77.0	9.1	6.14	1.38
VSR	30.4	12.0	4.06	2.59
DR-VDR Seq	77.0	9.2	6.12	1.37

Figure p. Mean and standard deviation of Active Share (AS) and Tracking Error (TE), computed weekly relative to the market-capitalisation-weighted benchmark. High standard deviation relative to mean indicates intermittent rather than structural active positioning.

Active weight distributions confirm the two-cluster structure identified in the active share analysis. VDR and DR-VDR Seq operate with approximately 75,000 active stock-date slots each (mean weights of 0.89% and 0.88% respectively), concentrated portfolios with P95 weights of approximately 3.3% and a maximum single-stock weight of 18.8% (VDR) and 17.2% (DR-VDR Seq). The slightly lower maximum weight for DR-VDR Seq reflects the DR anchor stage suppressing the most extreme single-stock concentrations that arise in the unconstrained VDR solution. DR and VSR operate very differently: approximately 210,000 active slots each with mean weights below 0.32%, consistent with broad, thin portfolios that blend passive market-cap exposure with marginal active tilts.

Figure r. presents time-averaged marginal contribution to risk (MCR) decomposition statistics for each strategy. DR-VDR Seq achieves the broadest risk distribution of any strategy: effective number of risk contributors (Eff-N) of 49.8, HHI of 0.0201, risk Gini of 0.807, and top-5 MCR share of 19.1%. This result is architecturally significant: the two-stage design, establishing a well-diversified DR anchor before imposing the VDR tilt, produces a more broadly diversified risk budget than standalone VDR (Eff-N 43.6), despite the two strategies running nearly identical weight books. The DR anchor stage's contribution is therefore primarily in risk space rather than return space: it prevents the VDR objective from concentrating the risk budget in a small number of high-signal names.

DR shows a risk decomposition (Eff-N 24.5, top-5 MCR 37.4%) that is weaker than its label implies. A strategy labelled by the Diversification Ratio objective should achieve the broadest risk distribution in the

group; instead it ranks below VDR and DR-VDR Seq on every MCR metric. This anomaly is consistent with the active share findings below and reflects the DR optimiser's convergence to a portfolio with high passive benchmark overlap rather than an aggressively diversified active book. VSR's risk decomposition (Eff-N 19.7, top-5 MCR 40.1%) reflects a near-passively-structured portfolio that concentrates residual active risk in a small number of positions, the opposite of what a volatility-scaling objective should produce.

Figure r. MCR Decomposition

Strategy	HHI (Risk)	Eff-N (Risk)	Risk Gini	Top-1 MCR %	Top-5 MCR %
Mcap	0.0802	12.5	0.955	21.5%	49.9%
DR	0.0408	24.5	0.916	10.3%	37.4%
VDR	0.0229	43.6	0.826	6.5%	21.9%
VSR	0.0509	19.7	0.917	13.7%	40.1%
DR-VDR Seq	**0.0201**	**49.8**	**0.807**	**4.2%**	**19.1%**



See Appendices for full size

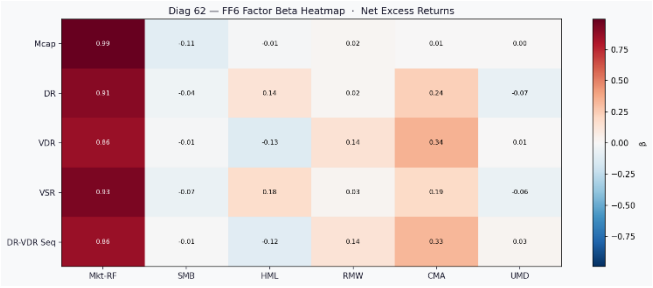
Marginal contribution to risk (MCR) decomposition statistics, time-averaged over the full backtest. HHI Risk is the Herfindahl-Hirschman Index of risk concentration. Eff-N is the effective number of risk contributors (1/HHI). Risk Gini measures inequality of risk distribution. Top-1 and Top-5 MCR are the percentage of total portfolio risk attributable to the largest single and five largest risk contributors respectively.

Figure s. presents full-period Fama-French six-factor (FF6) loadings. Several structural conclusions emerge. First, all optimised strategies reduce market beta relative to Mcap (0.994): DR 0.905, VDR 0.862, VSR 0.934, DR-VDR Seq 0.859. This beta reduction is the primary mechanical explanation for the drawdown protection documented above, and equivalently the primary return cost in high-beta bull markets.

Second, VDR and DR-VDR Seq share a distinctive factor profile: large negative HML (value) loading (−0.126, −0.116), meaningful RMW (profitability) loading (+0.141, +0.138), and CMA (conservative investment) loading (+0.344, +0.328). This combination, long profitable, low-investment, growth-oriented firms, is the factor-space signature of selecting stocks with stable, low volatility-of-volatility profiles. The RMW exposure (+0.14) distinguishes this pair from both DR and VSR and provides partial insulation from the CMA headwind: profitability was a positive factor in the 2023–2025 period even as conservative investment was a negative factor, explaining why VDR's 2023–2025 sub-period Sharpe (1.144) holds up better than DR-VDR Seq's (1.265, better, but with less RMW to cushion).

DR's factor profile has shifted materially relative to its design intention: an HML loading of +0.139 (positive value tilt) combined with a CMA loading of only +0.229 indicates the DR optimiser is selecting value names rather than low-investment quality names. This profile is inconsistent with a well-functioning diversification-ratio objective, which should produce a low-beta, low-investment factor tilt similar to VDR. The value tilt in DR is consistent with its active share finding, a near-passive portfolio that tilts modestly toward cheaper, larger names rather than pursuing genuine diversification. VSR's HML loading (+0.177) similarly indicates a value tilt, alongside the shallowest CMA exposure (+0.192) of any optimised strategy, confirming its near-passive character.

Figure s. FF6 Decomposition



See Appendices for full size

Fama-French six-factor model loadings, estimated via OLS regression over the full backtest period. Factors: Mkt-RF (excess market return), SMB (size), HML (value), RMW (profitability), CMA (investment), UMD (momentum).

A Brinson-Fachler decomposition was applied to isolate stock selection from sector/allocation effects relative to the full-universe total market benchmark. The results reveal a fundamental structural distinction between VDR and the three other optimised strategies. VDR is the sole strategy generating positive average selection alpha: +1.436% per annum (average weekly selection contribution of +0.000276, annualised). This confirms

that the volatility-of-volatility signal adds genuine stock-level discriminatory power beyond what can be attributed to sector weights or factor tilts, a finding consistent with the theoretical framing of VoV as a proxy for idiosyncratic risk regime changes not captured by first-order volatility measures.

DR, DR-VDR Seq, and VSR all exhibit negative selection contributions (DR −1.018%, DR-VDR Seq −0.723%, VSR −1.258% per annum), indicating that within-sector stock selection detracts from benchmark performance for each. The allocation drag is more modest than in the previous version of the backtest, DR −0.735%, DR-VDR Seq −0.261%, VDR −1.427%, VSR −0.410% per annum, reflecting that the new architecture confines the most aggressive active positioning to the VDR objective rather than distributing it across all strategies. DR-VDR Seq records the smallest allocation drag (−0.261%), a direct consequence of its DR anchor stage maintaining broader market coverage and limiting sector-level divergence from the benchmark.

VDR's interaction term warrants specific attention: at −0.936% per annum it is the largest in magnitude of any strategy, and it is negative, meaning that the sectors VDR overweights tend to be precisely the sectors where VDR's stock selection underperforms, and vice versa. This negative covariance between allocation and selection bets reflects the VoV signal operating in tension with the sector-level weight structure: VDR tilts toward sectors with lower aggregate VoV (allocation), yet within those sectors it selects the highest-VoV names it can find (selection), producing systematic cancellation at the interaction level. DR-VDR Seq's near-zero interaction term (+0.170%) confirms that the DR anchor stage resolves this tension by separating the sector-level diversification decision (stage one) from the stock-level VoV tilt (stage two), preventing the two objectives from working against each other simultaneously.

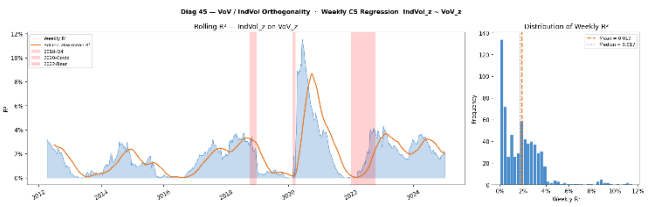
On a total Brinson basis, all four strategies produce negative annualised active return: DR −0.963%, VDR −0.927%, DR-VDR Seq −0.814%, VSR −1.549%. The active return shortfall is an unavoidable consequence of the volatility-based construction objective in the evaluation period: the highest market-capitalisation and benchmark-weight names, which the diversification and VoV objectives systematically underweight, were simultaneously the strongest return contributors during 2023–2024. DR-VDR Seq's smallest total active drag (−0.814%) confirms that the two-stage architecture minimises benchmark-relative return sacrifice while maintaining the structural risk advantages documented above.

e. VoV orthogonality & Blowout attribution

A central theoretical claim of the VDR framework is that volatility-of-volatility (VoV) constitutes an economically distinct information source from realised individual volatility. This claim was tested by regressing cross-sectional VoV scores on contemporaneous individual volatility scores weekly and recording the coefficient of determination (R^2). The results provide strong empirical support: the mean cross-sectional R^2 over the full backtest is 0.019 (median 0.017), and 95.5% of weeks exhibit R^2 below 0.05. VoV and realised individual volatility share less than 2% of common variance in a typical week.

The maximum observed R^2 is 0.116, occurring during the acute phase of the Covid-19 dislocation in March 2020, the one episode in which market-wide volatility explosion causes VoV and volatility levels to temporarily co-move. This spike is transitory; within 12 weeks of the March 2020 trough, R^2 returns below 0.04. The structural independence of the two signals across 95% of the sample period validates the VDR framework's core premise: strategies optimising on VoV capture diversification properties that are categorically unavailable to pure volatility-minimisation approaches.

Figure r. Volatility & VoV Orthogonality



See Appendices for full size

Figure s. examines weighted-average VoV exposure for each portfolio during the three principal stress episodes, reported as deviation from the contemporaneous equal-weighted (EW) universe average. A negative deviation would confirm that the VDR optimiser successfully underweights high-VoV names before and during a stress episode. Contrary to this prediction, all portfolios record positive deviations in all three events.

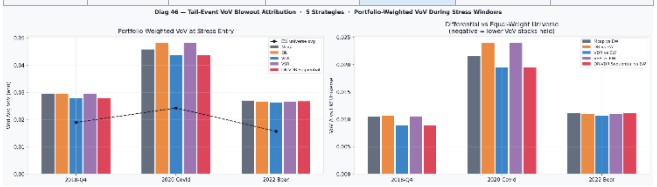
This result reflects a structural constraint rather than a failure of the VoV signal: the investable universe is restricted to the top-500 market-capitalisation segment, within which large-capitalisation stocks are subject to disproportionate institutional selling pressure, options market activity, and index-rebalancing flows during dislocations, all of which amplify short-term VoV

realisations relative to the broader equal-weighted universe. All portfolios sharing this universe are structurally exposed to this effect.

Within this constraint, VDR and DR-VDR Seq consistently achieve lower VoV blowout than both Mcap and DR. The ordering is clearest in 2020: VDR (+0.0194) and DR-VDR Seq (+0.0194) versus Mcap (+0.0215) and DR (+0.0239). In the 2022 bear market, all five portfolios converge toward similar VoV blowout levels (+0.0106 to +0.0112), reflecting the slower, less acute nature of the drawdown. The 2018-Q4 episode shows VDR and DR-VDR Seq at +0.0089 versus DR and VSR at +0.0106, a 16% reduction in excess VoV exposure. These findings are consistent with the VoV signal providing meaningful *relative* protection within the large-capitalisation universe, even if absolute protection against systematic volatility blowout is structurally unavailable.

Figure s. VoV Blowout Analysis

Event	EW VoV	Mcap vs EW	DR vs EW	VDR vs EW	VSR vs EW	DR-VDR vs EW
2018-Q4	0.0189	+0.0105	+0.0106	+0.0089	+0.0106	**+0.0089**
2020-Covid	0.0242	+0.0215	+0.0239	+0.0194	+0.0240	**+0.0194**
2022-Bear	0.0157	+0.0112	+0.0109	**+0.0106**	+0.0110	+0.0111



See Appendices for full size

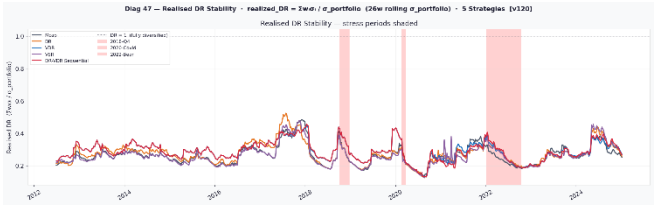
Weighted-average volatility-of-volatility (VoV) exposure for each portfolio during three stress events. Values reported as deviation from the equal-weighted (EW) universe average. Lower values indicate less VoV blowout relative to the average stock.

The realised Diversification Ratio was computed on a rolling basis as a measure of each strategy's achieved diversification relative to its objective. Over the full sample, DR-VDR Seq achieves the highest time-averaged realised DR (mean 0.302 across eras), followed by VDR (0.301) and DR (0.281). Mcap records 0.269, while VSR (0.264) falls below the benchmark, confirming that its near-equal-weight convergence produces less diversification benefit than the benchmark's passive construction.

The time-series of realised DR reveals a regime shift in 2023–2025: the gap between optimised strategies and Mcap compresses markedly, with all five portfolios converging toward similar realised DR values (0.283–0.292). This convergence reflects the rising within-index correlation driven by megacap concentration and passive ownership, as the 500-stock universe becomes more internally correlated, the marginal benefit of active diversification diminishes. The 2020–2022 stress period shows the largest divergence between the optimised strategies and Mcap, with VDR and DR-VDR Seq maintaining a 10-basis-point advantage in realised DR (0.258, 0.254 versus Mcap's 0.246).

An additional finding concerns VSR's 2023–2025 realised DR (0.292), which exceeds DR (0.290) and approaches DR-VDR Seq (0.290) in that period. This is an artefact of VSR's convergence to a near-equal-weight solution: the equal-weight portfolio achieves high realised DR as a consequence of equal weighting rather than optimisation. Realised DR therefore overstates VSR's active diversification contribution in recent periods and should be interpreted in conjunction with the active share and risk-decomposition diagnostics.

Figure t. Rolling Realised DR



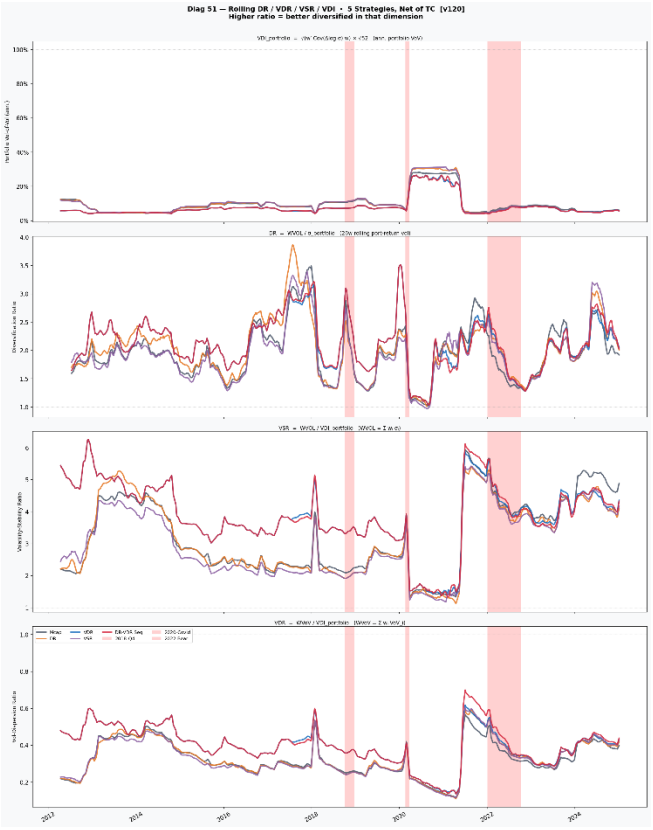
See Appendices for full size

Figure u. tracks the rolling realised values of all three ratio metrics (DR, VDR, VSR) for each portfolio, enabling cross-strategy comparison of achieved objective alignment. The full-sample means reveal a clear hierarchy. On the *VDR metric*, DR-VDR Seq achieves the highest realised value (0.4007), marginally above VDR (0.3962), confirming that the two-stage design does not sacrifice VDR optimisation relative to unconstrained VDR. Standalone DR achieves only 0.3350, confirming it is not a strong VDR portfolio. On the *DR metric*, DR-VDR Seq also leads (2.119), above VDR (2.110) and well above VSR (1.932), demonstrating that the first-stage DR maximisation in DR-VDR Seq is genuinely binding and produces a portfolio that dominates on both the DR and VDR dimensions simultaneously.

These cross-metric results confirm the theoretical motivation for the two-stage architecture: by anchoring to a genuine DR maximum in stage one, DR-VDR Seq achieves a higher realised DR than standalone VDR (+0.009, approximately +0.4%), while the VDR objective in stage two maintains or slightly improves the

achieved VDR relative to the unconstrained optimiser (+0.005). The DR anchor functions as a diversification floor that the VDR tilt cannot erode.

Figure u. Diversification metrics



See Appendices for full size

V) Conclusion

This paper introduced a Diversification Quality Framework extending the classical Diversification Ratio along a second dimension, the correlation of volatilities rather than the correlation of returns. Three new metrics were proposed to operationalise this dimension: the Volatility Dispersion Index (VDI), measuring how synchronously asset volatilities move; the Volatility Stability Ratio (VSR), measuring the stability of a portfolio's aggregate risk level relative to its average volatility; and the Volatility Diversification Ratio (VDR), diagnosing whether a portfolio's instability is driven by diversifiable idiosyncratic risk or by a common systematic risk factor. Together these metrics address the central limitation of the classical DR: a portfolio can appear well-diversified in return space while being entirely concentrated in volatility risk space, the condition identified here as False Systematic Risk.

The empirical results, evaluated over a 665-week backtest on the top-500 US equities from April 2012 to January 2025, confirm that this second dimension carries measurable and economically significant information. VDR and DR-VDR Seq, the two strategies directly optimising on the volatility-of-volatility signal, consistently outperform the benchmark and the classical DR optimiser on every risk-adjusted metric that matters institutionally: Sharpe ratio (0.962 and 0.951 versus Mcap's 0.923), Calmar ratio (0.546 and 0.543 versus 0.472), maximum drawdown (−25.22% and −25.23% versus −31.10%), CVaR at the 95th percentile (−231.85 and −235.08 bps versus −263.98 bps), and recovery speed from the March 2020 trough (16 weeks versus 20 for DR and Mcap and 34 for VSR). This performance advantage is not confined to isolated periods: DR-VDR Seq leads Mcap in 52.3% of 26-week rolling Sharpe windows, the only optimised strategy with a majority win rate across the full sample.

The cross-sectional R^2 analysis confirms the theoretical premise of the framework: VoV and realised individual volatility share less than 2% of common variance in a typical week, spiking only transiently to 11.6% at the acute phase of the Covid-19 dislocation before reverting below 4% within 12 weeks. Strategies optimising on VoV therefore access diversification properties that are categorically unavailable to pure volatility-minimisation approaches. This orthogonality is the empirical foundation on which the entire framework rests.

The DR-VDR Sequential architecture, the main portfolio construction contribution of this paper, demonstrates that the tension between DR and VDR, which can have mutually exclusive goals as identified in the framework, is resolvable through a two-stage SOCP rather than a scalar composite objective. By anchoring to a DR floor derived from a pre-computed high-DR solution in Stage 1, and then maximising the true VDR ratio subject to that floor via a second-order cone constraint in Stage 2, DR-VDR Seq simultaneously achieves the highest realised DR of any strategy (mean 0.302 versus VDR's 0.301 and DR's 0.281) and the highest realised VDR (0.4007 versus VDR's 0.3962). The DR floor does not sacrifice VDR optimisation, it improves it, by preventing the VDR objective from concentrating the risk budget into a small number of high-signal names. This is confirmed by the MCR decomposition: DR-VDR Seq achieves the broadest risk distribution of any strategy with an effective number of risk contributors of 49.8 versus VDR's 43.6 and DR's 24.5, despite running a near-identical weight book to VDR.

The stress-regime analysis provides the sharpest characterisation of each strategy's value proposition. During the 67 bottom-decile Mcap weeks, DR delivers the lowest mean loss (−2.82% per week), confirming that the classical diversification objective remains the most reliable capital protection mechanism in acute dislocations. DR-VDR Seq delivers the best tail-shape improvement, the least negative skewness of any strategy (−1.89 versus VDR's −2.38 and Mcap's −2.93) and the shallowest P5 loss (−6.14%), because the SOC DR-floor constraint specifically suppresses the extreme left-tail concentrations that the unconstrained VDR objective would otherwise allow. The combination captures both properties: mean protection from the DR floor, tail-shape improvement from the VDR objective operating within the SOC-constrained feasible region.

These results must be read honestly against two structural limitations. First, all optimised strategies underperform Mcap during the 2023–2025 megacap concentration phase, with DR-VDR Seq recording a mean calendar-year excess return of −1.24% over the full sample. This is not a failure of the framework but a direct consequence of the Portfolio Trilemma: a portfolio cannot simultaneously maximise return, minimise volatility, and minimise instability over long periods. The VoV framework explicitly trades concentration-driven return for risk-quality improvement, and the cost of that trade is paid most visibly in environments where a small number of megacap names dominate index return. Second, like all backward-looking optimisers, DR-VDR Seq allocates to the best protective assets based on past VoV behaviour. In the event of a genuinely new market paradigm, one where the historical VoV structure provides no signal about future co-movement, the framework offers no protection beyond what any diversification-based approach would provide.

VSR warrants specific mention as a negative result. Its Sharpe of 0.849, maximum drawdown worse than the benchmark, and annualised turnover of only 21.15%, lower than the passive Mcap benchmark, collectively indicate that the VSR optimiser converges to a near-equal-weight static solution rather than genuinely pursuing its stated objective. VSR's diagnostic value within the framework is real, as a measure of portfolio risk stability it correctly identifies under-diversified portfolios and market instability regimes, but as a standalone optimisation objective over a 500-stock long-only universe it is insufficient.

A useful conceptual parallel for the relationship between DR and VDR is the Delta-Gamma decomposition in options pricing. Traditional portfolio volatility acts as the portfolio's Delta, measuring its immediate first-order sensitivity to market moves. VDR, through VoV, acts as the portfolio's Gamma, quantifying the stability and predictability of that volatility estimate itself. A portfolio with reasonable volatility but low VDR is analogous to a Delta-neutral position with large unhedged Gamma: it appears controlled under static analysis but its risk profile is fundamentally unstable and can change violently with regime shifts. DR-VDR Sequential, by jointly constraining both dimensions, is the first systematic attempt to hedge both simultaneously within a single convex optimisation.

The Diversification Quality Framework does not replace classical portfolio theory, it completes it. Where Markowitz's efficient frontier maps the return-risk trade-off and the DR extends this to return-diversification, the VDR framework adds a third axis: the nature and stability of risk itself. By acknowledging that not all diversification is equal, and that the predictability and source of risk matter as much as its magnitude, the framework equips portfolio managers with the diagnostic tools to distinguish between portfolios that are genuinely diversified and those that merely appear so, and to build accordingly.

References

[1] Choueifaty, Yves, Towards maximum diversification (September 4, 2008). Available at SSRN: https://ssrn.com/abstract=4063676 or http://dx.doi.org/10.2139/ssrn.4063676

[2] Markowitz, Harry. "Portfolio Selection." *The Journal of Finance*, vol. 7, no. 1, 1952, pp. 77–91. *JSTOR*, https://doi.org/10.2307/2975974. Accessed 26 June 2026.

[3] Sharpe, William F. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *The Journal of Finance*, vol. 19, no. 3, 1964, pp. 425–42. *JSTOR*, https://doi.org/10.2307/2977928. Accessed 26 June 2026.

[4] Black, Fischer, and Myron Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, vol. 81, no. 3, 1973, pp. 637–54. *JSTOR*, http://www.jstor.org/stable/1831029. Accessed 26 June 2026.

[5] Merton, Robert C. "Theory of Rational Option Pricing." *The Bell Journal of Economics and Management Science*, vol. 4, no. 1, 1973, pp. 141–83. *JSTOR*, https://doi.org/10.2307/3003143. Accessed 26 June 2026.

[6] Engle, Robert F. "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica*, vol. 50, no. 4, 1982, pp. 987–1007. *JSTOR*, https://doi.org/10.2307/1912773. Accessed 26 June 2026.

[7] Fama, E.F. and French, K.R. (1992), The Cross-Section of Expected Stock Returns. The Journal of Finance, 47: 427-465. https://doi.org/10.1111/j.1540-6261.1992.tb04398.x

[8] Dan Rosen, David Saunders, Risk factor contributions in portfolio credit risk models, Journal of Banking & Finance,Volume 34, Issue 2, 2010, Pages 336-349, ISSN 0378-4266, https://doi.org/10.1016/j.jbankfin.2009.08.002. (https://www.sciencedirect.com/science/article/pii/S0378426609001964)

[9] Alexander, C., & Korovilas, D. (2011). The hazards of volatility diversification. ICMA Centre Discussion Paper in Finance No. DP2011-04. Available at SSRN: https://ssrn.com/abstract=1752389

[10] Stambaugh, Robert F. and Yu, Jianfeng and Yuan, Yu, Arbitrage Asymmetry and the Idiosyncratic Volatility Puzzle (October 12, 2014). Journal of Finance, Vol 70, pp 1903-1948, 2015, Jacobs Levy Equity Management Center for Quantitative Financial Research Paper , Available at SSRN: https://ssrn.com/abstract=2155491 or http://dx.doi.org/10.2139/ssrn.2155491

[11] Hsu, Stephen D. H. and Murray, Brian, On the Volatility of Volatility (August 24, 2006). Available at SSRN: https://ssrn.com/abstract=926379 or http://dx.doi.org/10.2139/ssrn.926379

[12] Baltussen, Guido and van Bekkum, Sjoerd and van der Grient, Bart, Unknown Unknowns: Uncertainty About Risk and Stock Returns (October 23, 2014). Journal of Financial and Quantitative Analysis, 53(4), p.1615-1651, 2018. , Journal of Financial and Quantitative Analysis, 53(4), p.1615-1651, 2018. , Available at SSRN: https://ssrn.com/abstract=2023066 or http://dx.doi.org/10.2139/ssrn.2023066

[13] Schlag, Christian and Shaliastovich, Ivan and Thimme, Julian and Huang, Darien, Volatility-of-Volatility Risk (May 2018). SAFE Working Paper No. 210, Jacobs Levy Equity Management Center for Quantitative Financial Research Paper, Available at SSRN: https://ssrn.com/abstract=3183610 or http://dx.doi.org/10.2139/ssrn.3183610

[14] Branger, Nicole and Hülsbusch, Hendrik and Kraftschik, Alexander, The Volatility-of-Volatility Term Structure (January 9, 2018). Paris December 2017 Finance Meeting EUROFIDAI - AFFI, Available at SSRN: https://ssrn.com/abstract=2980074. or http://dx.doi.org/10.2139/ssrn.2980074

[15] Ilmanen, Antti S. and Kizer, Jared, The Death of Diversification Has Been Greatly Exaggerated (2012). Journal of Portfolio Management, Vol. 38, No. 3, pp. 15-27, 2012, Available at SSRN: https://ssrn.com/abstract=2998754

[16] Page, S., & Taborsky, M. A. (2011). The myth of diversification: Risk factors versus asset classes. The Journal of Portfolio Management, 37(4), 1-2

[17] Roncalli, Thierry and Weisang, Guillaume, Risk Parity Portfolios with Risk Factors (September 21, 2012). Available at SSRN: https://ssrn.com/abstract=2155159 or http://dx.doi.org/10.2139/ssrn.2155159

[18] Hjalmarsson, Erik, Diversification Across Characteristics (December 2009). FRB International

Finance Discussion Paper No. 986, Available at SSRN: https://ssrn.com/abstract=1536765 or http://dx.doi.org/10.2139/ssrn.1536765

[19] Shumway, T. (1997), The Delisting Bias in CRSP Data. The Journal of Finance, 52: 327-340. https://doi.org/10.1111/j.1540-6261.1997.tb03818.x

[20] Beaver, William H. and McNichols, Maureen F. and Price, Richard A., Delisting Returns and Their Effect on Accounting-Based Market Anomalies. Journal of Accounting and Economics (JAE) 43, issues 2+3, July 2007, 341-368, Available at SSRN

Appendices

Full-scale figures

Figure d. Proposed Fully Diversified Portfolio Framework

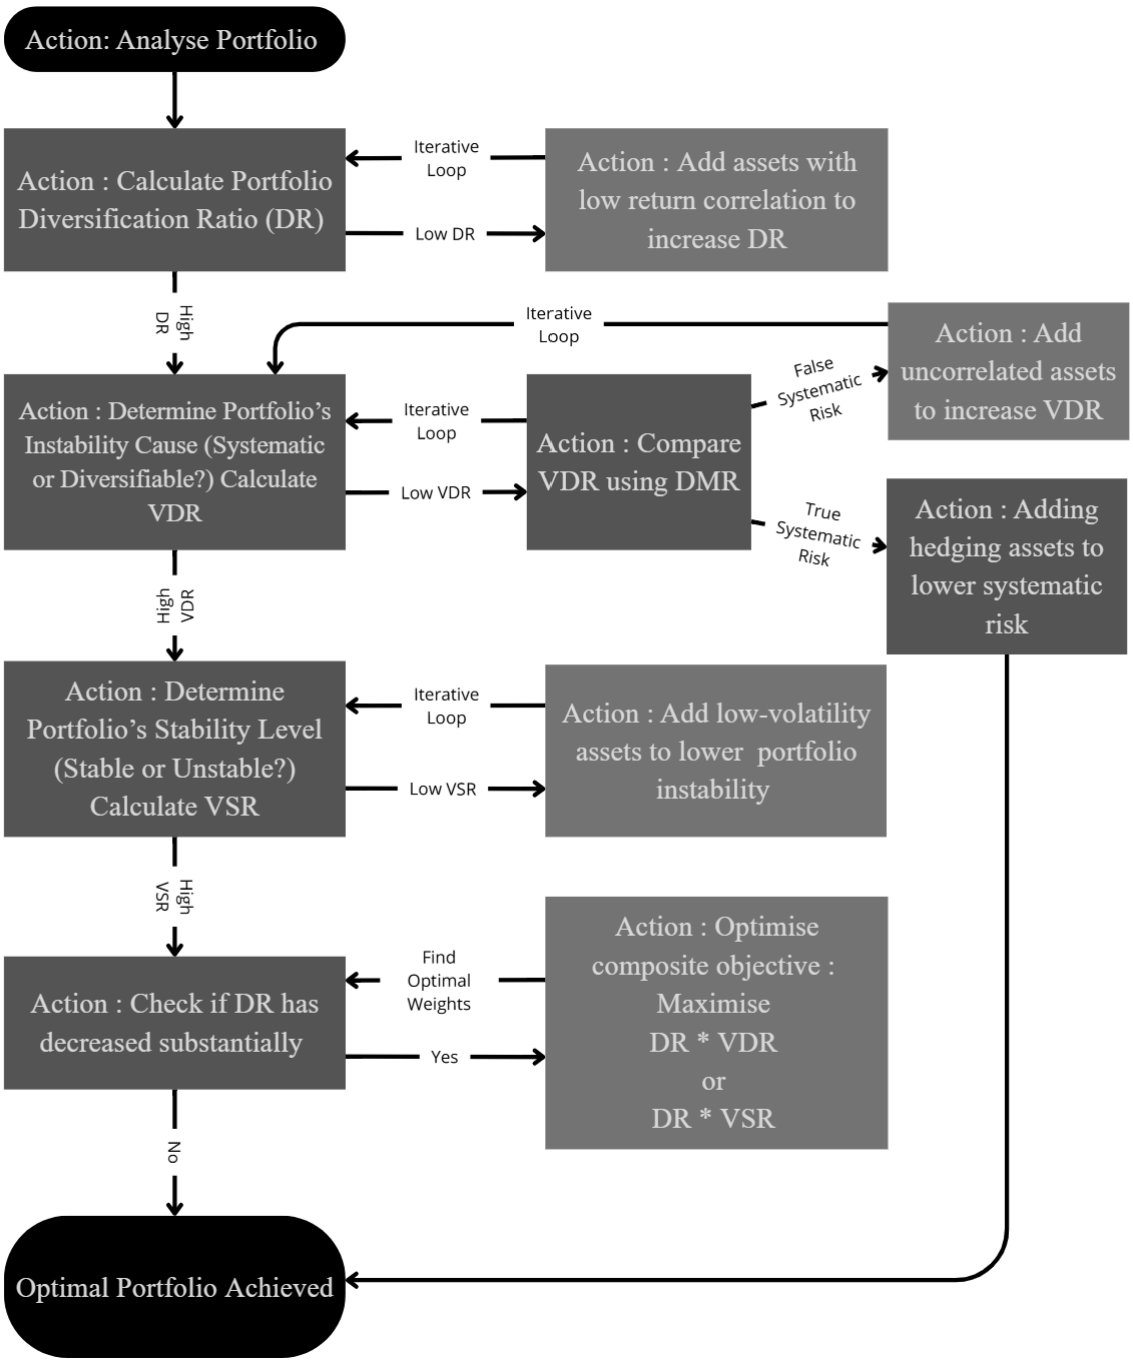


Figure f. Portfolio's Volatility-Stability Efficient Frontier

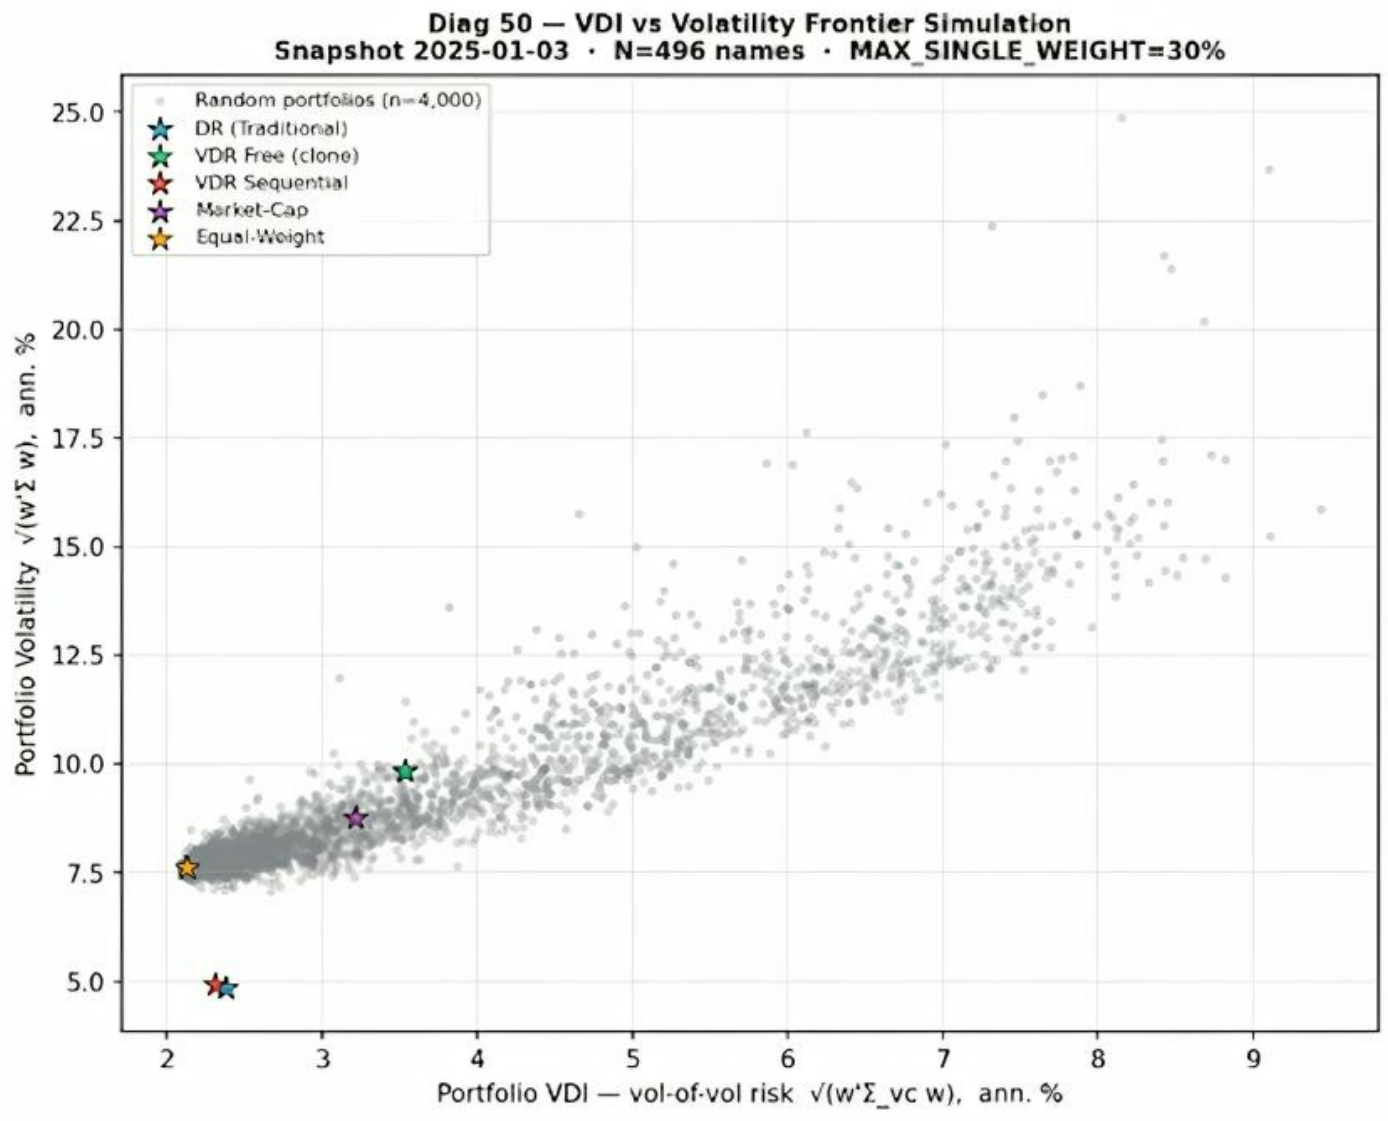


Figure g. Portfolio's Volatility-Stability Efficient Frontier + Return

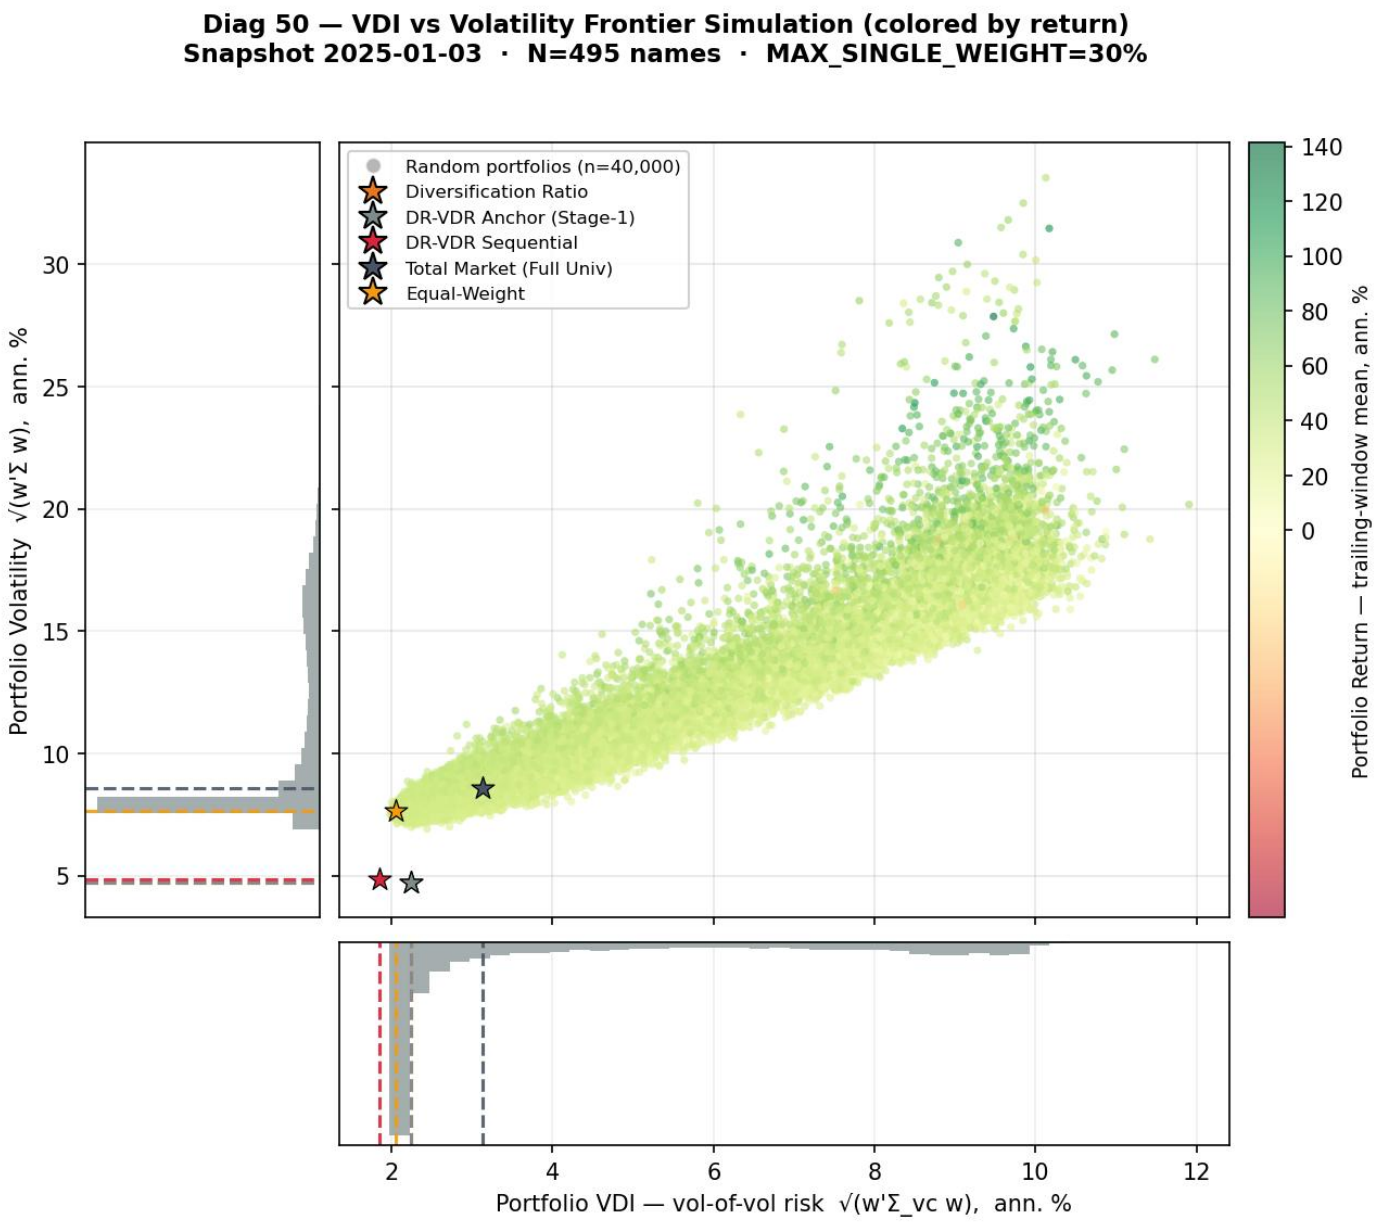


Figure h. Portfolios' Realised DR

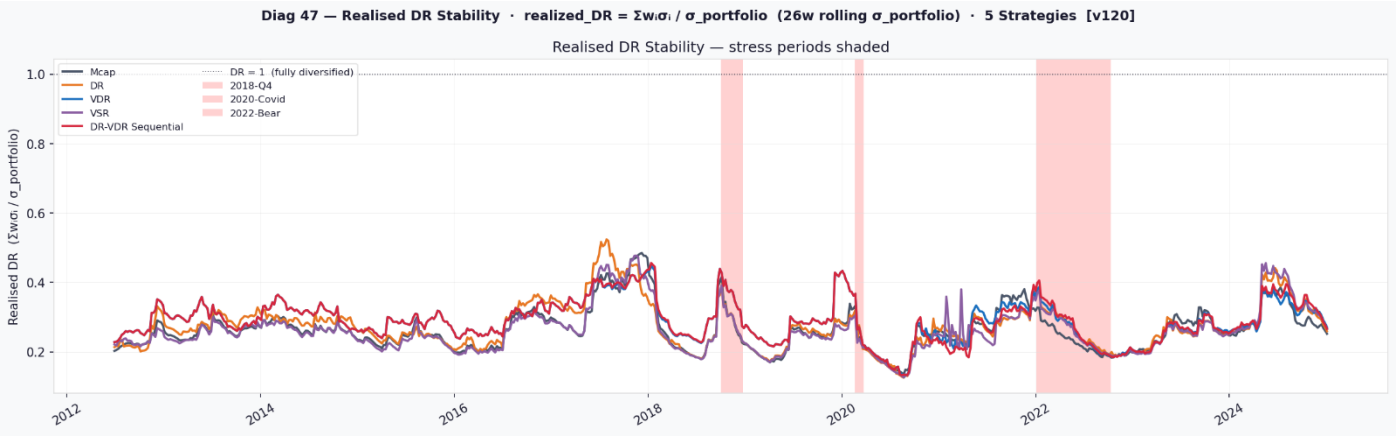


Figure i. Portfolios steps

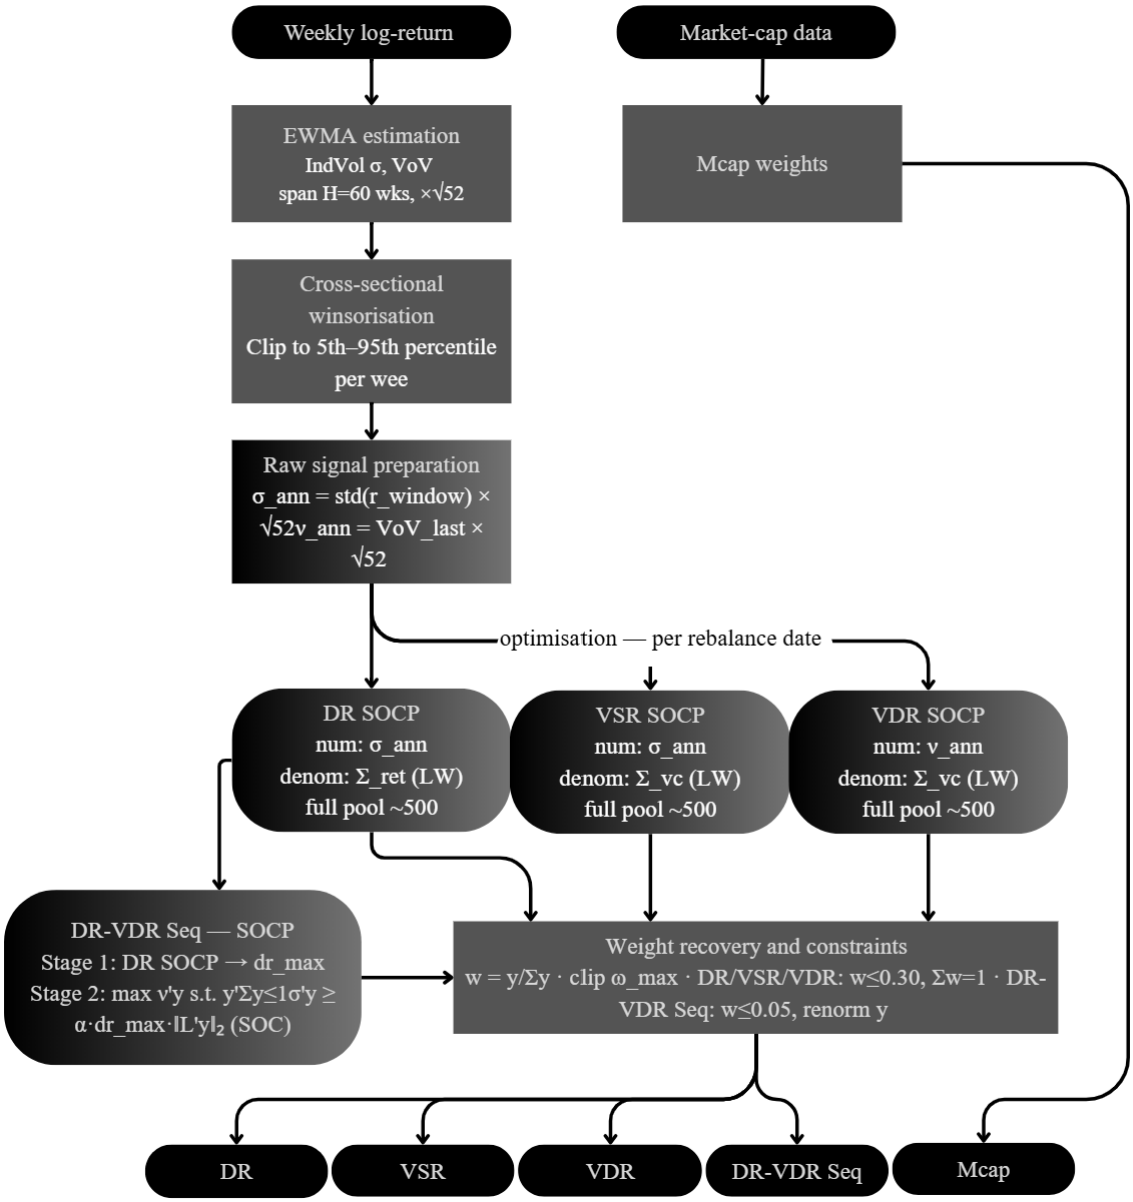


Figure k. Portfolio NAV evolution

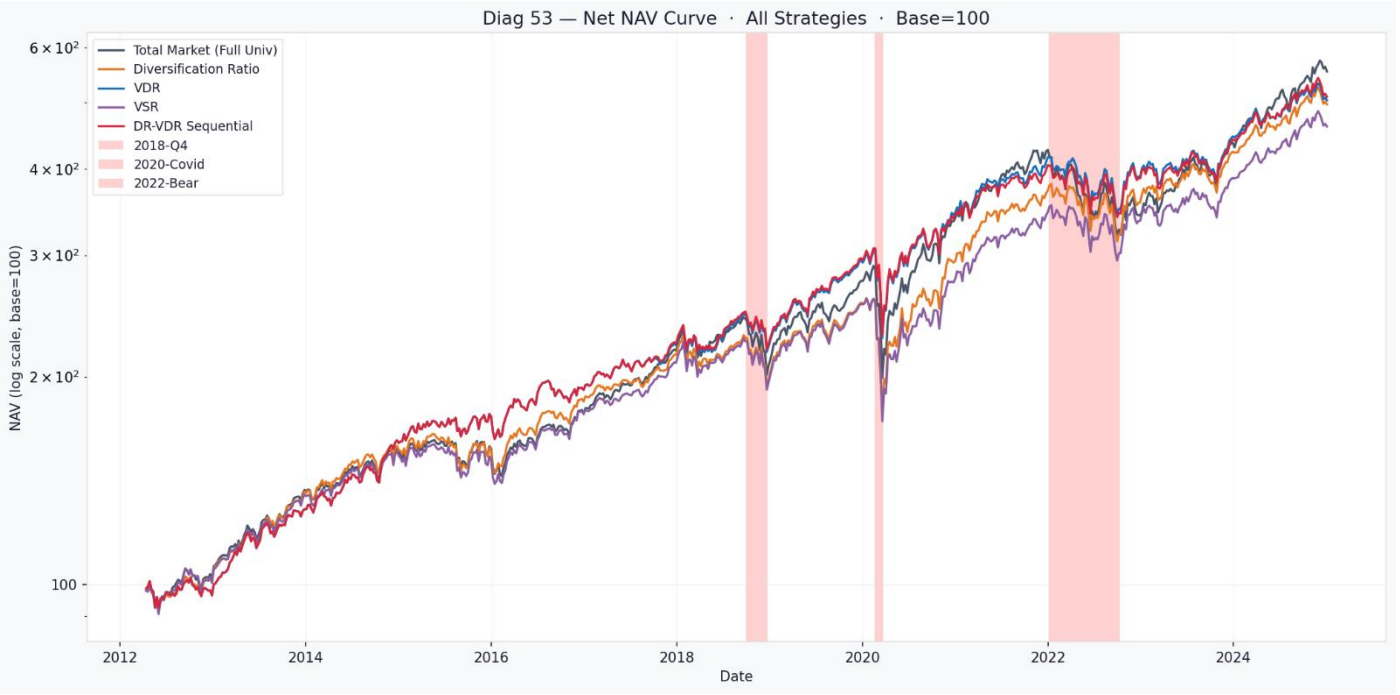


Figure l. Sub-Period Sharpe Analysis

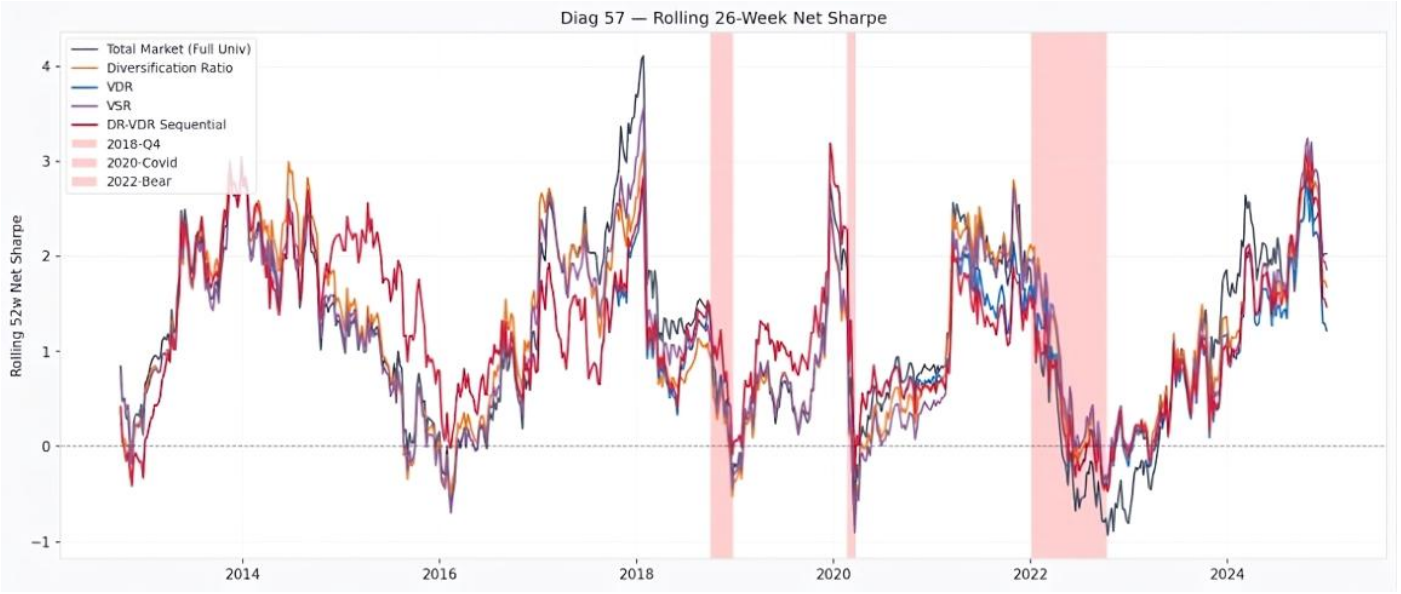


Figure m. Calendar Year Returns

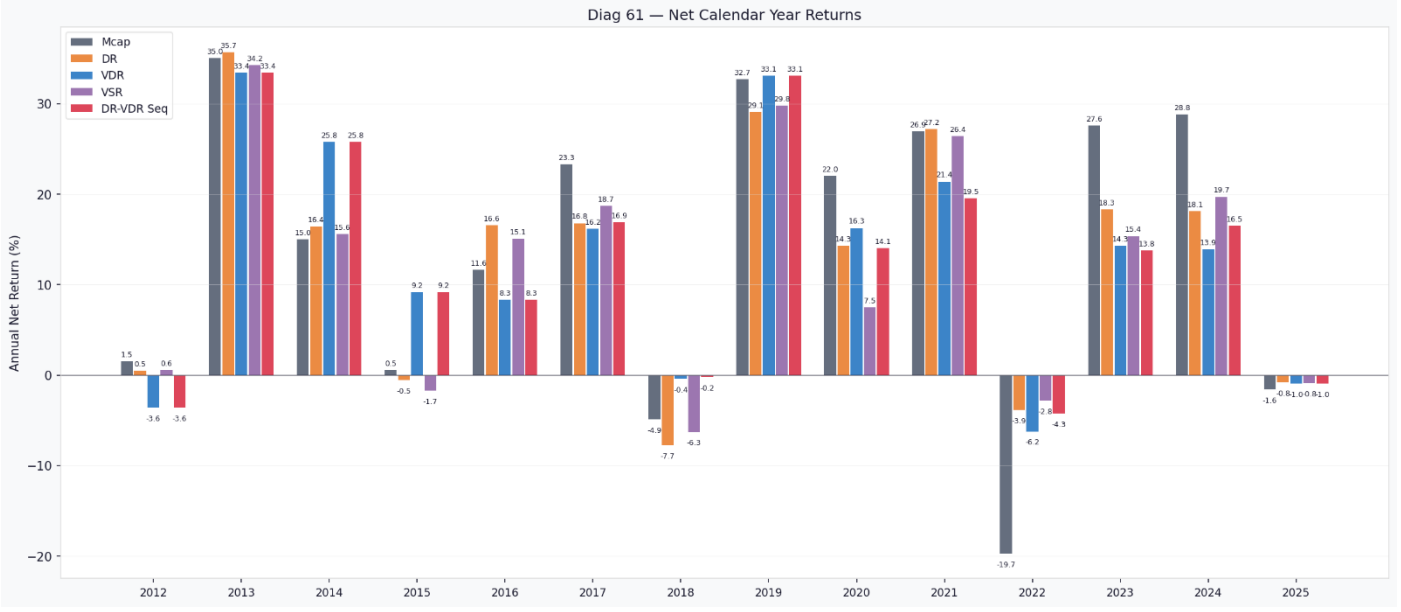


Figure n. Conditional Regime Distribution

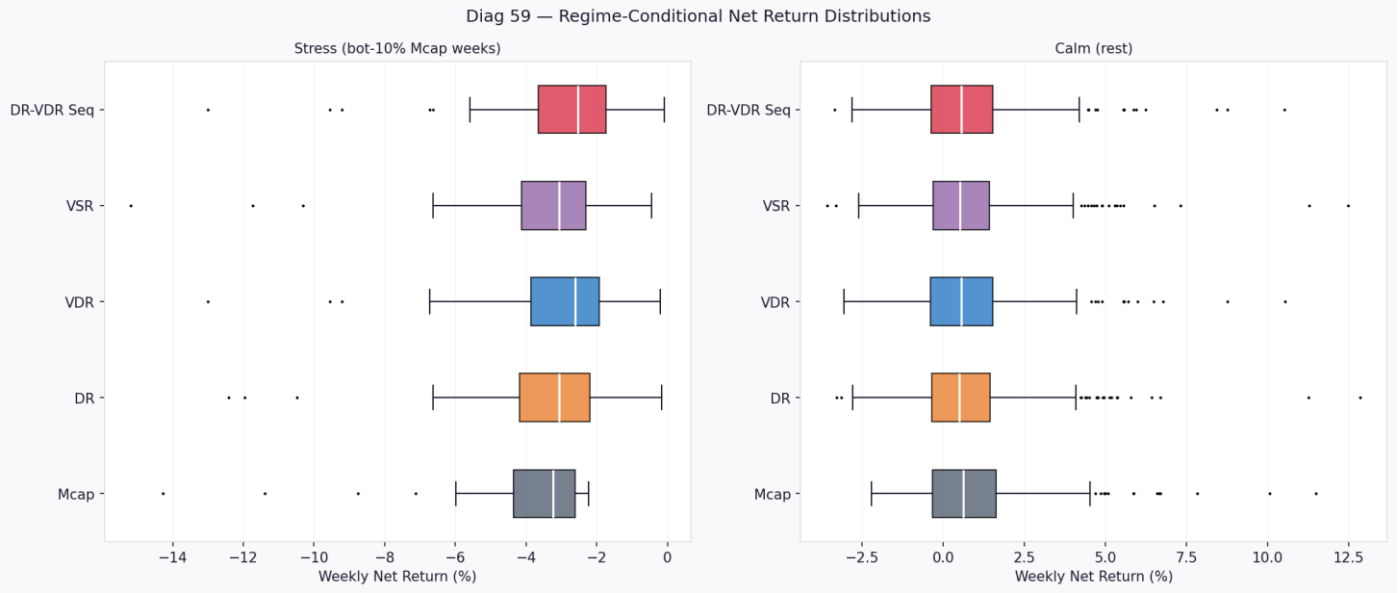


Figure o. Return Asymmetry

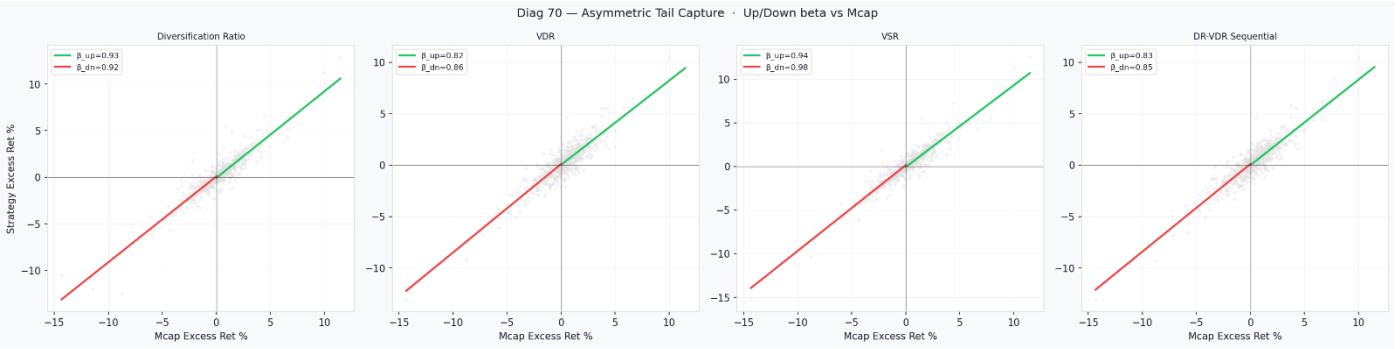


Figure p. Drawdown analysis

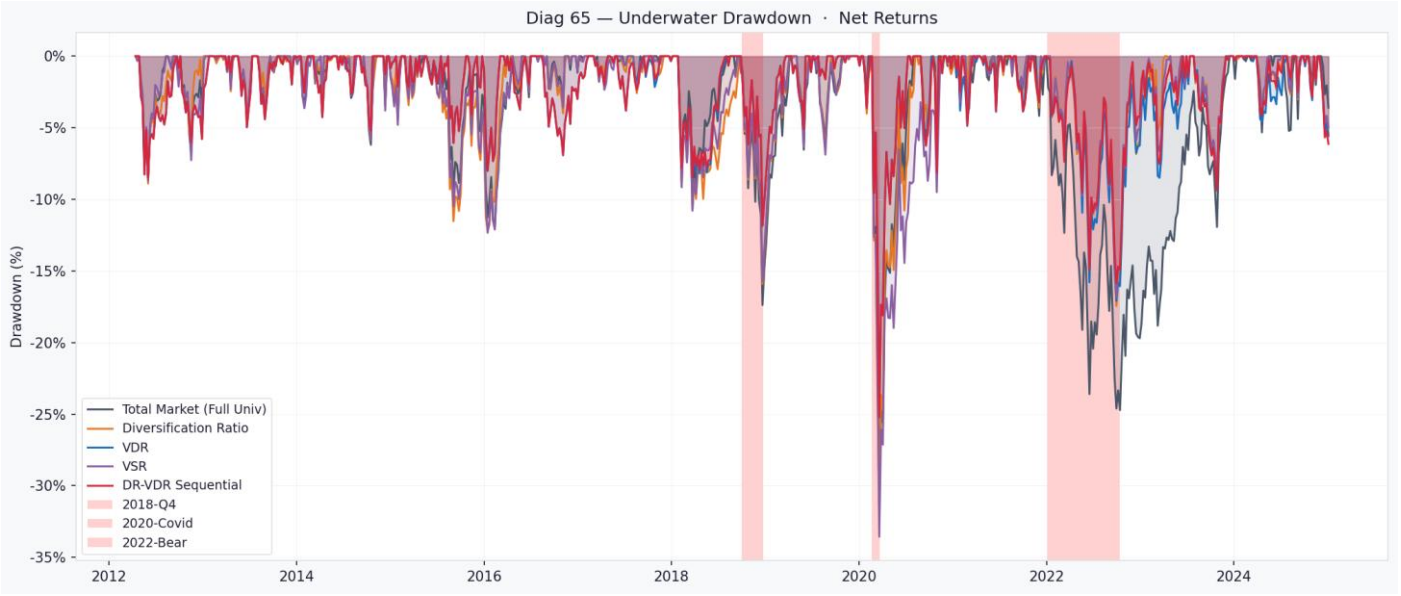


Figure q. MCR Decomposition



Figure r. FF6 Decomposition

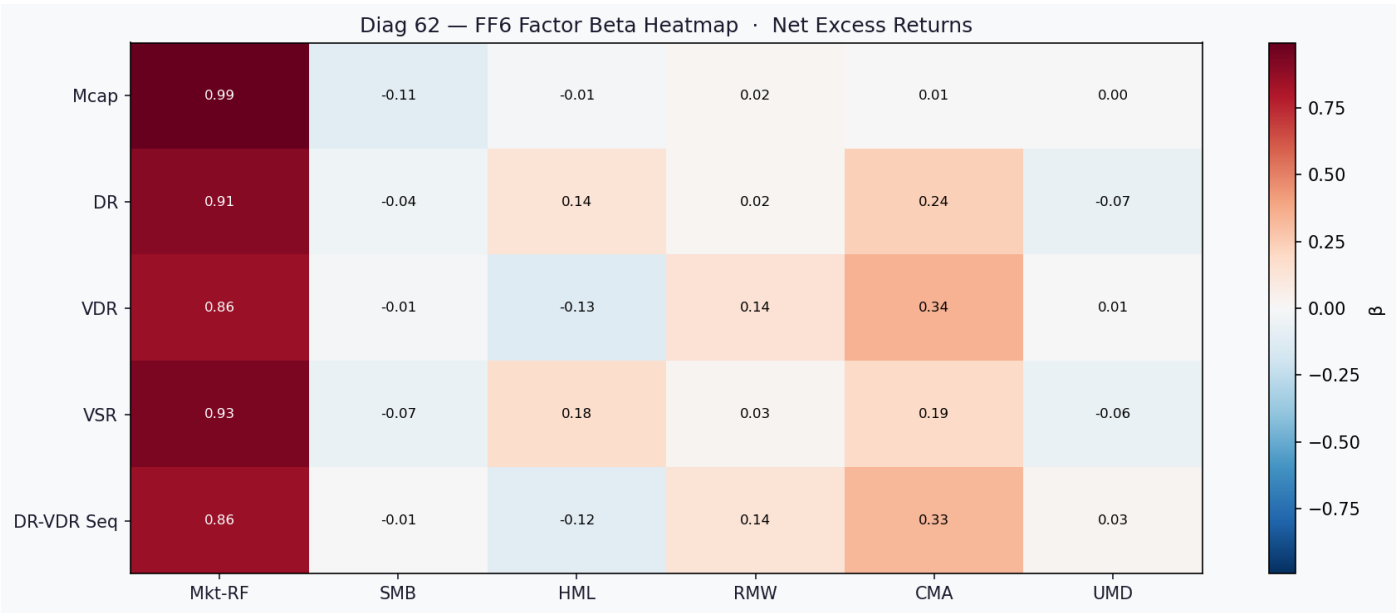


Figure s. Volatility & VoV Orthogonality

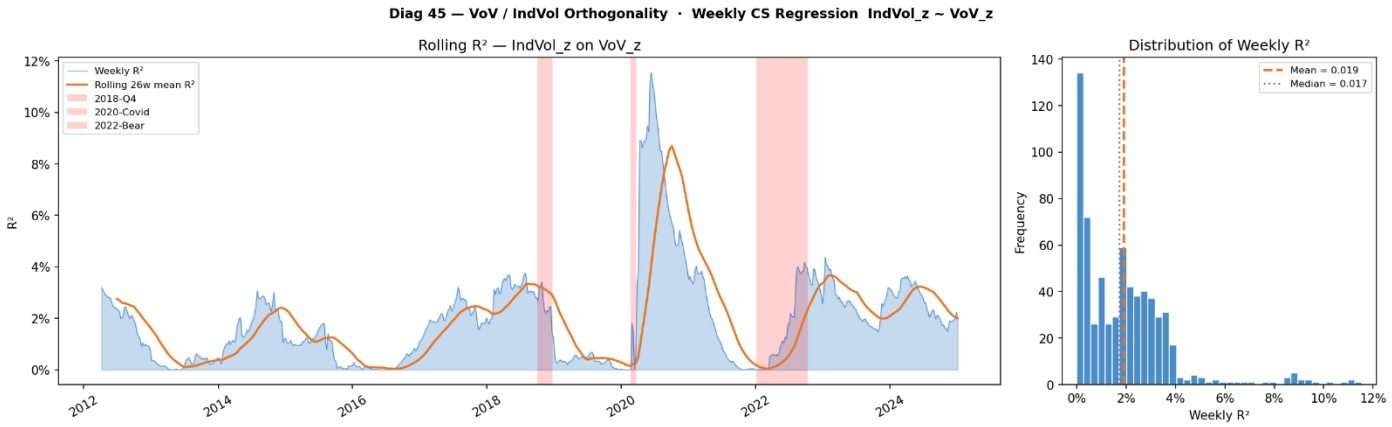
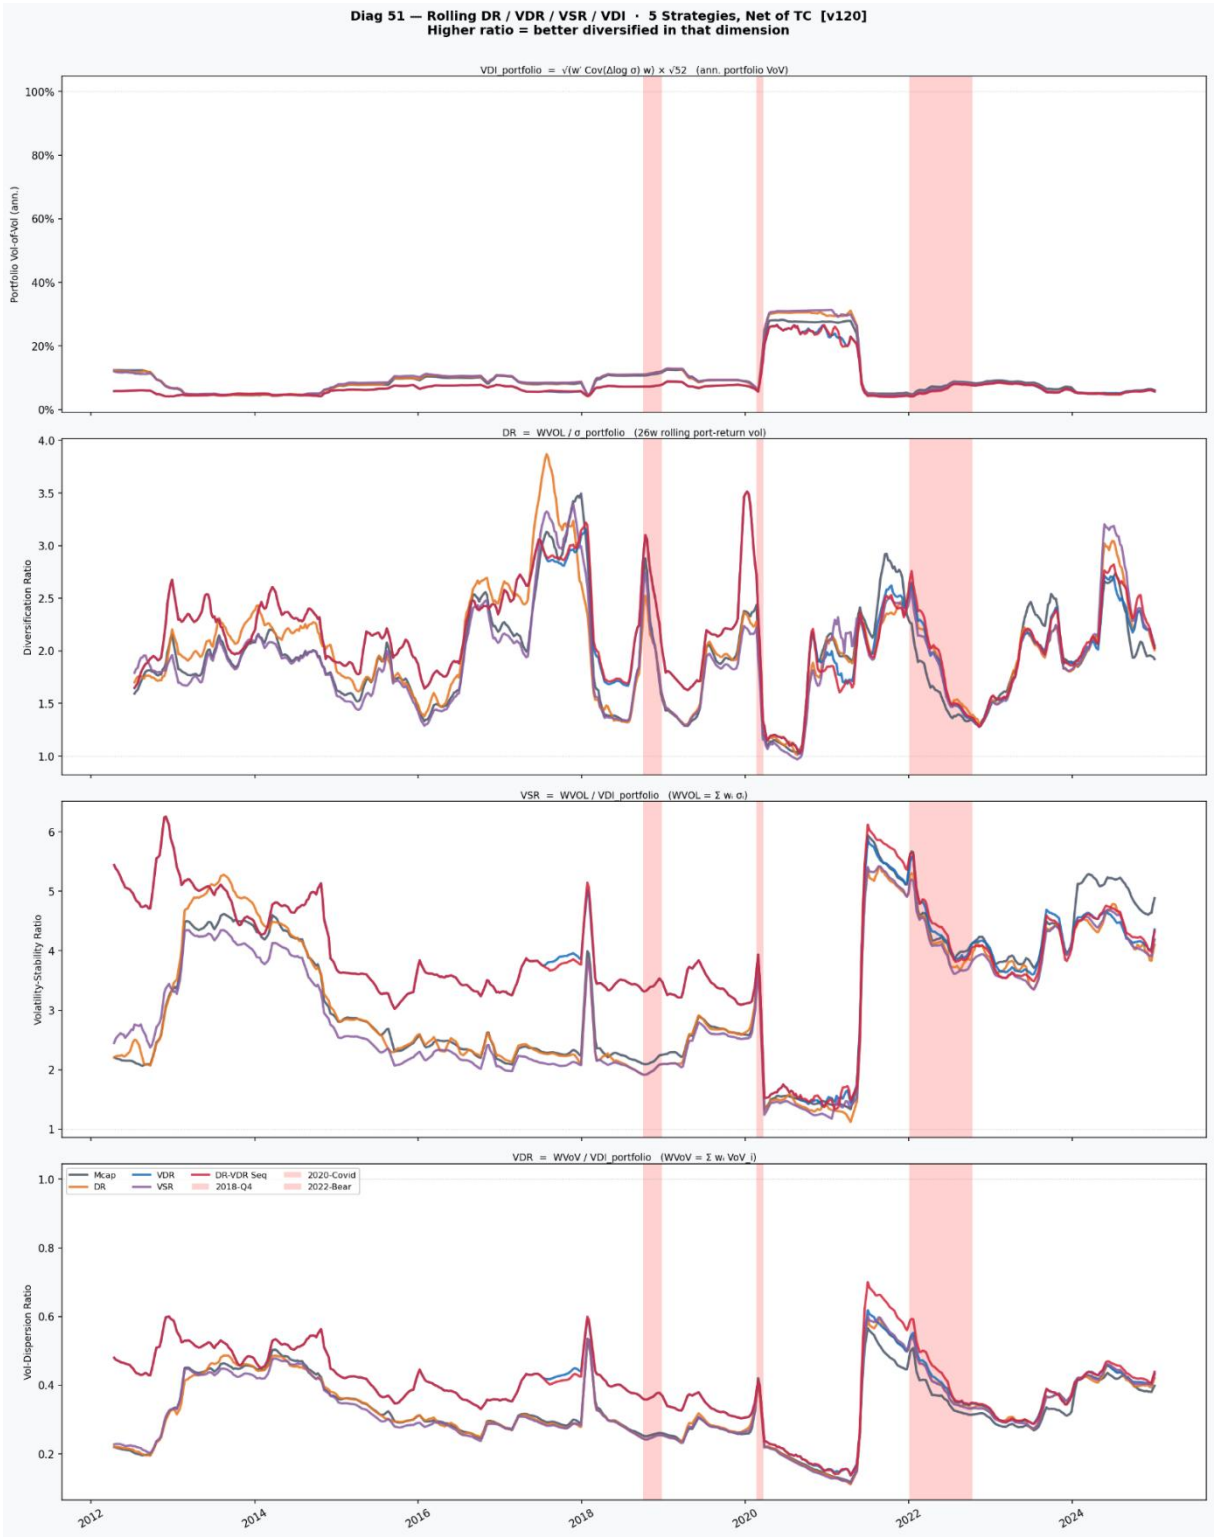


Figure t. VoV Blowout Analysis



Figure u. Diversification metrics


Diag 51 — Rolling DR / VDR / VSR / VDI · 5 Strategies, Net of TC [v120]
Higher ratio = better diversified in that dimension
VDI_portfolio = √(w' Cov(Δlog σ) w) × √52 (ann. portfolio VoV)
Portfolio Vol-of-Vol (ann.)
100%
80%
60%
40%
20%
0%
DR = WVOL / σ_portfolio (26w rolling port-return vol)
Diversification Ratio
4.0
3.5
3.0
2.5
2.0
1.5
1.0
VSR = WVOL / VDI_portfolio (WVOL = Σ wᵢ σᵢ)
Volatility-Stability Ratio
6
5
4
3
2
1
VDR = WVoV / VDI_portfolio (WVoV = Σ wᵢ VoV_i)
Vol-Dispersion Ratio
1.0
0.8
0.6
0.4
0.2
Mcap
DR
VDR
VSR
DR-VDR Seq
2018-Q4
2020-Covid
2022-Bear
2012
2014
2016
2018
2020
2022
2024