

Historical Covariance Forecasts Underestimate Stress Risk by a Factor of 1.6

We tested four ways to build a portfolio using 54 years of daily US industry data. Every month we check the risk forecast against the volatility that actually arrives in the month that follows. The models consistently underestimate risk during market stress.

SUBJECT

Covariance estimation error and correlation in market stress

EVIDENCE

49 US industries daily 1972-2026, five regions, US Treasuries

DOCUMENT TYPE

Research report

In the 33 most volatile months since 1973 every covariance model's median risk forecast was too low, and the textbook optimiser underestimated risk in all 33

The covariance matrix is fitted on a past that is much calmer than the moment it is needed. Pure estimation error makes the optimiser's forecast optimistic even in quiet markets, while during shocks all tested allocation methods miss reality by a similar magnitude. Closing the gap requires a faster volatility estimate paired with a stressed correlation ^[01].

TEXTBOOK
OPTIMISER STRESS
FAILURES

100%

Share of volatile months where the model underestimated risk.

OPTIMISER STRESS
FORECAST MISS

1.68

Median ratio of realised to predicted volatility in market shocks.

EQUAL WEIGHT
CALM BIAS

0.74

Median ratio showing how the naive allocation overstates risk.

STRESSED
EXPONENTIAL
MODEL

0.99


Realised volatility ratio using fast decay and stressed correlation.



OBSERVATION	A mean-variance model requires an impossibly long estimation window to beat a naive equal-weight allocation out of sample ^[05] . Yet during market stress, the equal-weight and hierarchical portfolios underestimate risk just as severely as the optimisers ^[01] . The gap stems from the covariance inputs rather than the portfolio math.
INTERPRETATION	Turbulence shreds the historical covariance structure exactly when you rely on it most. You can shrink the matrix or build a complex hierarchy, but neither fixes inputs that are too calm.
IMPLICATION	Risk teams must decouple the volatility estimate from the correlation estimate. Accelerate the volatility decay and hardcode a stressed correlation assumption to survive market shocks.
CAVEAT	We define stress months retrospectively by their realised volatility, mechanically selecting periods where any trailing forecast will look too low. The test also isolates US equity industries. This ignores the cross-asset dynamics that might cushion a real portfolio shock.

In a typical year only 2 of 49 eigenvalues rise above the noise band

With 49 assets and one year of days, a random matrix already produces eigenvalues up to 2.08^[06]. The market factor carries more than half the variance. The rest cannot be mathematically distinguished from noise, yet the optimiser inverts all of it.



EIGENVALUES ABOVE NOISE EDGE

2 / 49

Median count out of 49 across 643 rolling windows.

VARIANCE IN FIRST EIGENVALUE

58.2%

Median share of total variance captured by the market factor.

CONDITION NUMBER

536

Median ratio of the largest to smallest eigenvalue across all windows.

INTERPRETATION

Standard optimisers treat every matrix component as hard fact. They place massive bets on the smallest eigenvalues to force theoretical diversification. Matrix inversion mathematically amplifies this error, letting random static dictate the final portfolio weights.

IMPLICATION


Risk teams must strip the noise before optimisation. Stop feeding raw empirical covariance matrices into your models and apply spectral filtering or shrinkage instead.

CAVEAT

The theoretical noise band assumes independent and identically distributed returns. Weak but genuine economic relationships can easily hide inside it. Small eigenvalues do not necessarily lack information, they just lack the strength to survive matrix inversion.

With one year of data the optimiser understates its own risk by a quarter even when nothing changes

This bias needs no crisis. A mean-variance optimiser simply loads on the directions that look least risky by pure chance ^[03]. Ledoit-Wolf shrinkage ^[04] roughly halves this bias, to 1.13 in our simulation, but does not remove it ^[01].



UNDERSTATEMENT FACTOR

1.24

With one year of data.

GROSS POSITION

367%

Median in the walk-forward test.

SHORT POSITIONS

133%

Median in the textbook model.

OBSERVATION

Testing 14 optimising models reveals none consistently beat a naive equal-weight allocation out of sample ^[05]. An unconstrained mean-variance model requires an estimation window of 6,000 months for 50 assets to reliably outperform equal weights ^[05].

INTERPRETATION


The algorithm treats sample noise as structural truth. It finds asset combinations with low historical variance and builds massive offsetting positions. The resulting portfolio reflects the errors in your specific sample rather than the actual market.

CAVEAT

The simulation assumes normal returns and a stable covariance. This isolates pure estimation error and understates the real problem. Real optimisers usually add position limits, which cut both the extreme leverage and part of the bias.

Every risk model proves too optimistic in a crisis, and the naive allocations too cautious in calm markets

Equal-weighted and hierarchical risk parity ^[10] portfolios routinely overstate risk in quiet markets. Their realised volatility runs at roughly 0.75 of the forecast. Both models then violently understate risk in a crisis, hitting a median ratio near 1.65. The textbook optimiser fares worse: it understates risk even in calm months (1.10). Shrinkage is accurate in calm months (0.99) but still misses by 1.63 in stress.




PORTFOLIO	CALM MONTHS	ALL MONTHS	STRESS MONTHS	STRESS MONTHS UNDERESTIMATED
Textbook optimiser	1.10	1.20	1.68	100.0%
Shrinkage optimiser	0.99	1.08	1.63	93.9%
Hierarchical risk parity	0.75	0.86	1.65	84.8%
Equal weights	0.74	0.85	1.63	87.9%

INTERPRETATION	Trailing correlations break down exactly when you need them to hold. You build a covariance matrix on a year of quiet data, but volatility and correlations both gap higher in a shock. The risk estimate fails mainly because volatility jumps; rising correlation adds the rest.
CAVEAT	These portfolios hold entirely different weights. Their absolute realised volatilities naturally diverge. We isolate the forecasting error here. The ratio compares a specific model's risk prediction only against its own eventual outcome.

In the worst tenth of months industry correlation reaches 0.62, against 0.45 in ordinary up months

Risk models build their forecasts with the correlations of the trailing year. During the stress episodes shown below, the month itself was far more correlated than the recent past ^[01]. The overall curve is asymmetric: the best months also push correlation higher, just not as violently ^[07].




EPISODE	CORRELATION, TRAILING YEAR	CORRELATION, THAT MONTH	REALISED / FORECAST RISK (EQUAL WEIGHT)	MARKET, MONTH
October 1987	0.56	0.86	6.12	-22.6%
August 1998	0.51	0.72	2.00	-15.7%
October 2008	0.62	0.83	3.22	-17.2%
August 2011	0.65	0.91	3.26	-5.9%
March 2020	0.52	0.83	6.70	-13.2%
April 2025	0.37	0.75	3.23	-0.4%

OBSERVATION	In stress months the trailing year averages a correlation of 0.58. The stress month itself jumps to 0.71. Calm months do the opposite and drop from 0.53 to 0.44 ^[01] .
INTERPRETATION	Diversification vanishes exactly when investors need it most. An equal-weight portfolio expects a diversification ratio of 1.38 heading into a stress month. The realised ratio falls to 1.20 ^[01] .
CAVEAT	Correlation moves toward one, it does not reach it. The highest 63-day average on record is 0.86, set in August 2011. Correlation measured in a volatile month is partly inflated by the volatility itself ^[09] .

Beyond two standard deviations down, correlation is 0.75 where a normal model expects 0.27

A normal model says extreme days should be less correlated. The data say the opposite. A covariance matrix, which describes the middle of the distribution, cannot carry this ^[07] ^[08].




OBSERVED DOWNSIDE CORRELATION 0.75 Average across industries when both asset and market fall past two standard deviations.	NORMAL BENCHMARK 0.27 Theoretical expectation for the exact same extreme downside moves.	OBSERVED UPSIDE CORRELATION 0.65 Average across industries when both asset and market rise past two standard deviations.
---	---	---

INTERPRETATION	Covariance is a linear metric built on calm days. It assumes extreme moves will scatter. The data show they converge. A risk model relying on this metric is blind to the actual mechanics of a market crash.
IMPLICATION	Portfolios need stress tests built on empirical tail dependence. Diversification models fed by historical covariance will fail precisely when the market breaks.
CAVEAT	The downside to upside gap in our data is smaller than the gap to the normal model. Tail dependence drives this result, alongside asymmetry. Forbes and Rigobon show volatility regimes mechanically inflate measured correlation ^[09] .

Rising correlation explains a fifth of the stress miss; rising volatility the rest

The prevailing narrative claims that correlations go to one in a crisis. Our decomposition of the equal-weight portfolio shows something else: most of the variance miss comes directly from each industry becoming more volatile ^[01]. Rising correlation multiplies the miss by a further 1.26, about a fifth of the total in log terms. Yet that specific fraction happens to be the cushion diversification was supposed to provide when you need it most.



CORRELATION SHARE OF MISS

20.3%

Share of the log variance prediction error, leaving the rest to rising volatility.

STRESS DIVERSIFICATION RATIO

1.20

Falls from 1.59 in calm months, indicating a severe loss of portfolio cushioning.

REALISED STRESS CORRELATION


0.71

Average pairwise coefficient during the shock, rising from 0.58 in the trailing forecast window.

OBSERVATION	Over 33 stress months the equal-weight portfolio generated a variance 3.09 times higher than predicted ^[01] . Re-pricing the forecast with the actual stress volatilities but the stale trailing correlations explains a factor of 2.45. The change in correlation accounts for the remaining multiple of 1.26, or 20.3% of the total log variance miss.
INTERPRETATION	Volatility does the heavy lifting in breaking your risk budget. Individual industries gap down, and their standard deviations explode. The correlation breakdown simply removes the remaining diversification benefit exactly when the portfolio takes the brunt of the volatility shock.
IMPLICATION	Risk teams obsess over modelling complex correlation dynamics when they should first stress-test baseline volatility. If you cannot survive a pure volatility shock, a better correlation matrix will not save you.
CAVEAT	This exact split depends on the mathematical order of the decomposition. Computing the correlation effect first would yield slightly different attribution shares. The analysis also relies specifically on an equal-weight allocation, so a heavily concentrated portfolio might show a different balance between the two effects.

The 99% VaR broke 2.25% of days, and ten of 52 years ended in the Basel red zone

Normal models expect failure to be rare and random. In the data, breaks cluster violently in stress years. When a breach happens, the resulting loss towers over what a normal distribution allows ^[01].



DAILY BREAK RATE

2.25%

against the 1% failure rate expected under normality

YEARS IN THE RED ZONE

10/52

of 52 full calendar years measured by Basel rules

SEVERITY MULTIPLE

1.47

average loss on a break day as a multiple of VaR, compared to a normal 1.15

OBSERVATION	Measured in forecast standard deviations, the empirical 97.5% expected shortfall hits 3.31. A normal model expects 2.34. The extremes scale even worse: the worst break on record destroyed 9.0 times the estimated VaR in October 1987.
INTERPRETATION	Trailing covariance maps the center of the distribution. Risk management happens in the tails. By the time a portfolio hits a stress event, the matrix has already stopped describing reality, leaving capital buffers sized for a calm market that no longer exists.
IMPLICATION	Risk teams rely on unadjusted normal VaR at their own peril. Internal limits require expected shortfall metrics calibrated directly to historical stress.
CAVEAT	This specific test applies a normal VaR to a static equal-weight portfolio. Banks typically rely on historical-simulation VaR.

Faster volatility plus a stressed correlation brought the stress-month forecast to 0.99, at a cost in calm months

Speed matters more than the correlation assumption. A stressed correlation alone helps less than an exponentially weighted covariance with the RiskMetrics decay of 0.94 ^[11]. Combined, they match the median stress month, but they still break Value at Risk on 6.3% of stress days. Overstating risk in calm months is the inevitable price.




MODEL	STRESS, REALISED / FORECAST	STRESS MONTHS UNDERESTIMATED	STRESS DAYS BREAKING VAR	CALM, REALISED / FORECAST
A · Trailing one-year covariance	1.65	87.5%	11.98%	0.74
B · Exponentially weighted covariance	1.21	62.5%	8.08%	0.83
C · Trailing volatility, stressed correlation	1.41	78.1%	9.13%	0.61
D · Exponential volatility, stressed correlation	0.99	50.0%	6.29%	0.67

INTERPRETATION	Reacting quickly to rising volatility prevents more damage than assuming extreme correlation. Models fail in stress primarily because variance explodes. Fixing the correlation matrix without speeding up the volatility estimate leaves the system too slow.
IMPLICATION	Prioritise short-memory volatility estimators over static correlation buffers. You must accept inflated risk estimates during calm periods to cover the actual risk of a crisis.
CAVEAT	Estimating correlations only on the worst tenth of market days is a failed variant. It lowered the forecast by 9.7% because truncated samples bias measured correlation downward ^[07] . Model D's stressed correlation is still taken from the past.

The stock - bond correlation that protected portfolios from 2000 to 2021 has turned positive in 62.5% of months since 2022.

Two decades of negative correlation provided the window almost every multi - asset covariance matrix in use was fitted on. The 1970- 1999 record was positive. On the worst stock days since 2022, Treasuries no longer rallied reliably ^[01] ^[02].



CORRELATION 2000 - 2021

-0.37

The negative correlation era that trained modern risk models.

AUGUST 2026 CORRELATION

+0.30

The highest reading since June 2024.

REGIONAL EQUITY CORRELATION

0.76


Average across five regions since 2015, up from 0.46 in the 1990s.

OBSERVATION	On the worst 1% of stock days from 2000 to 2021, the 10 -year yield fell 9.7 basis points on average and dropped on 83.6% of those days ^[01] ^[02] . Since 2022, the yield rose 2.2 basis points on average and fell on only 41.7% of the worst equity days ^[01] ^[02] .
INTERPRETATION	A positive correlation regime means stocks and bonds sell off together, stripping multi -asset portfolios of their built -in shock absorber. Risk models calibrated purely to the previous two decades blindly expect bonds to save the day.
CAVEAT	The post -2022 sample of worst equity days contains only 12 days. Yield changes provide a mechanical proxy for bond price direction, not a precise return series.

Industry correlation is at 0.17, the third lowest reading since 1972, which is when forecasts are most optimistic

A covariance matrix estimated today carries one of the calmest correlation structures on record ^[01]. April 2025 proved how fast it can double. History shows these extreme lows reliably give way to much higher readings.

63-DAY AVERAGE PAIRWISE CORRELATION



63-day average pairwise correlation of daily returns, 49 US industries, month ends, 2000 to August 2026 ^[01].

AUGUST 2026 CORRELATION

0.17

The third lowest reading of 654 month ends since 1972.

REALISED RISK MULTIPLE

3.23

Equal-weight portfolio forecast mismatch in April 2025.

HISTORICAL REVERSION RATE

77.8%

Share of lowest-tenth readings followed by a correlation jump above 0.50 within two years.

INTERPRETATION

Models trained on current data generate massive position sizes because they perceive immense diversification benefits. During April 2025 the market dropped just 0.4%, yet correlation jumped to 0.75 and the textbook optimiser realised 4.06 times its predicted volatility.

IMPLICATION

A portfolio team does not need to wait for the rule. The European Commission postponed its application of the Basel 97.5% expected shortfall standard calibrated to stress ^[12] to 1 January 2027 ^[14]. Compute risk with stressed correlations today.

Historical covariance models consistently underestimate risk during market stress

We tested standard portfolio construction methods over decades of data. In the 33 most volatile months, every historical covariance model's median forecast was too low, and the textbook optimiser underestimated risk in all 33. The underlying math breaks down exactly when investors need protection.

- 01 A trailing year of data leaves just 2 of 49 eigenvalues above the noise band ^[06]. The optimiser understates true volatility by 1.24 times on this sample.
- 02 All four methods' median stress ratio was 1.63–1.68.
- 03 The change in correlation explains only 20.3% of the variance miss, leaving rising volatility to account for the rest.
- 04 The standard normal Value at Risk model broke its limits enough to trigger the Basel red zone in 10 of 52 years ^[13].
- 05 Applying exponentially weighted volatilities to a stressed correlation matrix aligns the forecast with realised risk. This combination achieves a 0.99 ratio during market shocks. The model degrades during calm months, where realised volatility is only 0.67 of the forecast.

SUMMARY OF FORECAST ERRORS

MEASURE	WHAT IT SHOWS	VALUE
Optimiser stress misses	Share of stress periods where the textbook model underestimated risk	33 / 33
Optimiser stress ratio	Realised volatility divided by forecast in the most volatile months	1.68
Correlation miss share	Portion of the variance forecast error driven by changing correlations	20.3%
VaR break rate	Frequency of daily portfolio losses exceeding the modelled risk limit	2.25%
Model D stress ratio	Realised against forecast volatility using stressed correlation and exponential weights	0.99
August 2026 industry correlation	Average trailing correlation of industries at the end of August	0.17

Before taking a risk number to an investment committee, pin down three inputs. Ask for the estimation window, the speed of the volatility estimate, and the exact correlation assumed in stress.

Selected sources

-
- [01] **Kenneth French Data Library**
https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html
 We extracted daily returns for 49 industry portfolios, the US market and regional markets, matching the CRSP and Bloomberg August 2026 builds.
-
- [02] **FRED, 10-year Treasury constant maturity rate (DGS10)**
<https://fred.stlouisfed.org/series/DGS10>
 We used this series for daily bond yields to track stock-bond correlation.
-
- [03] **Michaud, The Markowitz Optimization Enigma: Is 'Optimized' Optimal?, Financial Analysts Journal 1989**
<https://www.newfrontieradvisors.com/media/t04a3oct/the-markowitz-optimization-enigma.pdf>
 Michaud recognised early on that mean-variance optimisers act as estimation-error maximisers.
-
- [04] **Ledoit and Wolf, A well-conditioned estimator for large-dimensional covariance matrices, Journal of Multivariate Analysis 2004**
https://econpapers.repec.org/article/eejmvana/v_3a88_3ay_3a2004_3ai_3a2_3ap_3a365-411.htm
 We implemented their shrinkage estimator to test whether dampening sample noise improves risk forecasts.
-
- [05] **DeMiguel, Garlappi and Uppal, Optimal Versus Naive Diversification, Review of Financial Studies 2009**
https://econpapers.repec.org/article/ouprfnst/v_3a22_3ay_3a2009_3ai_3a5_3ap_3a1915-1953.htm
 The authors showed that an optimising model needs 6,000 months of data to beat naive diversification for 50 assets.
-
- [06] **Laloux, Cizeau, Potters and Bouchaud, Random matrix theory and financial correlations, 2000**
<https://www.cfm.com/wp-content/uploads/2022/12/234-1999-random-matrix-theory-and-financial-correlations.pdf>
 Their work defines the noise band we used to measure the random share of eigenvalues.
-
- [07] **Longin and Solnik, Extreme Correlation of International Equity Markets, Journal of Finance 2001**
https://econpapers.repec.org/article/blajfinan/v_3a56_3ay_3a2001_3ai_3a2_3ap_3a649-676.htm
 Exceedance correlation lets us track dependence deep in the tails.
-
- [08] **Ang and Chen, Asymmetric correlations of equity portfolios, Journal of Financial Economics 2002**
https://econpapers.repec.org/article/eejfinec/v_3a63_3ay_3a2002_3ai_3a3_3ap_3a443-494.htm
 The data confirm downside correlations run 11.6% above normal distribution estimates.
-
- [09] **Forbes and Rigobon, No Contagion, Only Interdependence, Journal of Finance 2002**
https://econpapers.repec.org/article/blajfinan/v_3a57_3ay_3a2002_3ai_3a5_3ap_3a2223-2261.htm
 They proved that higher volatility mechanically biases correlation upward.
-
- [10] **López de Prado, Building Diversified Portfolios that Outperform Out of Sample, Journal of Portfolio Management 2016**
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2708678
 We ran hierarchical risk parity to see if bypassing full covariance inversion helps out of sample.
-
- [11] **J.P. Morgan and Reuters, RiskMetrics Technical Document, 1996**
<https://www.msci.com/research-and-insights/paper/1996-riskmetrics-technical-document>
 We built an exponentially weighted covariance benchmark with a decay factor of 0.94.
-
- [12] **Basel Committee, MAR33 Internal models approach: capital requirements calculation**
<https://www.bis.org/committees/bcbs/basel-framework/standard/mar/33/inforce/2023-01-01/published/2020-06-05>
 This standard shifts trading-book capital requirements to a 97.5% expected shortfall.
-
- [13] **Basel Committee, MAR99 Guidance on use of the internal models approach**
<https://www.bis.org/committees/bcbs/basel-framework/81331/chapter.pdf>
 The backtesting rules put a bank in the yellow zone from 5 exceptions and the red zone from 10, per 250 trading days.
-
- [14] **European Commission, press release of 12 June 2025**
https://finance.ec.europa.eu/news/commission-proposes-postpone-one-additional-year-market-risk-prudential-requirements-under-basel-iii-2025-06-12_en
 The Commission delayed the EU application of these market-risk rules to 1 January 2027.
-

This report provides research material rather than investment advice. We calculated all results on historical US equity industries before transaction costs. Past estimation error does not establish a hard limit on future model failure.

Independent research, introductions clearly defined.

BlackRidge is an independent research bureau introducing private investors to quantitative traders through multiple strategy providers. We publish quantitative research. Investors access strategies through PAMM accounts at the broker. We are not a fund or broker and never hold client money.

01

Your funds stay at the broker. The account is opened in your own name. BlackRidge never receives, holds, or has withdrawal rights over client capital.

02

No fee on your profits. A one-time access payment of 6.7% of the agreed trading level, paid to the selected strategy provider. No management fee, no performance fee, no profit share, in any year.

03

You fund the risk, not the exposure. Allocations are notionally funded: you agree a trading level and fund the margin and drawdown allowance behind it. Trading losses can exceed the deposit without applicable negative balance protection; the separate 6.7% access payment is nonrefundable.

CONTINUE READING

blckridge.com/research; published research. [Notional funding](#) and [PAMM accounts](#) explained in full, and [answers](#) to the questions this raises.

We use AI models to gather and aggregate source material and to help prepare each report. Read it with its cited sources, sample, methods and limitations, and send corrections through [research methodology and corrections](#).

The client pays the selected strategy provider a one-time access payment of 6.7% of the agreed trading level. BlackRidge receives an introduction fee from that provider and does not receive broker compensation. [About BlackRidge](#).

Minimum funded capital \$25,000. Strategy authorship, live records, the broker, the selected strategy provider and the full commercial terms are presented on an introductory call and confirmed in writing before any payment or deposit.

This report is published for information only. It is not investment advice and does not take account of your circumstances. Trading leveraged instruments including CFDs carries substantial risk and is not suitable for all investors; you may lose the capital you fund and, depending on your broker's terms, may owe more than your deposit. Findings in this report may combine third-party evidence, historical calculations and illustrative scenarios; they are not verified live trading results of any strategy provider. Past performance does not indicate future results.