

A Formula from 1952, Solved Every Morning

*Ten mean-variance portfolios over 192 stock
tokens, with an active-set solver implemented in
Stylus*

ISOQUANT

1 October 2026

Abstract

In 1952 Harry Markowitz showed how to choose a portfolio by trading expected return against variance. We implement this calculation for a universe of 192 stock tokens using a single-factor risk model and ten fixed levels of risk aversion. A transaction on a local Nitro node running Stylus used 1,226,229 gas for a synthetic-data estimation and optimization benchmark, or 3.8% of the 32 million gas transaction limit read from Robinhood Chain. This measurement excludes the additional cost of reading the historical return tape from persistent storage, publishing price data, and executing portfolio trades. The return model assigns each stock the bill rate plus a fixed premium for its estimated market exposure; it does not extrapolate each stock's historical return. The optimizer identifies an active set and solves its optimality equations analytically. Its fixed-point implementation is evaluated against numerical optimality conditions and a double-precision reference. On prices through 30 September 2026, point 5 holds approximately 99% SPY. This concentration is consistent with a model that uses SPY as its market factor; point 5 is a fixed risk-aversion portfolio, not a permanently defined tangency portfolio. A one-year walk-forward simulation reports returns from 3.8% to 162.3% before trading costs, against 15.2% for SPY. The results demonstrate computational feasibility under the stated model and benchmark conditions. They do not establish live execution performance or the safety of the proposed vault design.

Keywords mean-variance optimization, single-index model, tangency portfolio, on-chain computation, Stylus

JEL classification C61, G11, G12

1 Introduction

Markowitz (1952) formulated portfolio selection as a trade-off between expected return and covariance. Tobin (1958) added a riskless asset and described how investors can combine it with a tangency portfolio. Sharpe (1964) connected the market portfolio to that equilibrium framework. These

ideas remain useful foundations for portfolio construction, but applying them requires choices about inputs, constraints, and execution.

This paper asks a narrower engineering question: can a blockchain compute a useful family of mean-variance portfolios at a modest transaction cost? We use Sharpe’s single-index structure to reduce the covariance model to a diagonal component and one market factor. The resulting optimality conditions allow an active-set method with an analytical closing step, avoiding a long sequence of projected-gradient iterations.

For a universe of 192 stock tokens, the reported local Stylus benchmark computes ten portfolios using 1,226,229 gas. The benchmark measures estimation and optimization; persistent price-history reads and portfolio execution add costs. Automated advisers also perform services beyond this calculation, so the benchmark is not a like-for-like comparison with an adviser’s management fee. Our contribution is the implementation and evaluation of this structured computation in a constrained execution environment. Figure 1 shows the model portfolios on the last day of the dataset.




Figure 1. The board on 30 September 2026. Grey dots are the 185 stocks with risk below 120%; the solid curve is the long-only frontier through the ten points, dotted where it runs past point 10. The dashed blue reference line runs from the bill rate through point 5; tangency is not certified.

Section 2 describes the inputs and risk model. Section 3 derives the optimality conditions and explains the implemented solver. Section 4 reports what it costs, Section 5 what the ten portfolios hold today, and Section 6 what they would have done over the past year. Section 7 describes the proposed vault design.

2 The Board and the Model

2.1 The board

The board is the set of 192 stock tokens on Robinhood Chain on 30 September 2026. It includes three bond funds (SGOV, SHY and BND), broad index funds such as SPY, VTI and QQQ, and single companies from Apple to small quantum-computing firms. A stock joins the morning solve once it has sixty daily returns; on 30 September, 191 of the 192 had them.

2.2 Risk

For each stock the solve regresses its last sixty daily returns on those of SPY. This gives the stock's market exposure β_i and the variance of what is left over, d_i , its own noise. With f the market's variance, all annualized, the covariance of the board is

$$\Sigma = D + f \beta \beta^\top, \quad (1)$$

where D is diagonal with entries d_i . This is the single-index model of Sharpe (1963). It needs $2n + 1$ numbers where a full covariance matrix needs $n(n + 1)/2$, which for this board is 385 against 18,528.

The implementation floors each annualized residual variance at 2^{-20} , approximately 9.54×10^{-7} , equivalent to about 0.098% annual residual volatility. In the Rust Q40 representation this is stored as 2^{20} ; the JavaScript model uses the same real-valued floor. This makes D positive definite, including for SPY, whose regression on itself otherwise has zero residual variance. The floor is a regularization choice, not measured residual risk; its sensitivity has not been reported.

2.3 Expected returns

The solve also needs an expected return for each stock, and here it declines to guess. Past returns estimate future ones badly, and a mean-variance optimizer magnifies the error by putting the most money where the estimate is most wrong (Michaud, 1989). The year to 30 September made the point plainly. AXT rose 1,483% and Sandisk 1,433%, while NuScale Power fell 79%. An optimizer that believed those numbers would hold little besides last year's winners. We set instead

$$\mu_i = r_f + \beta_i \pi, \quad (2)$$

where r_f is the annualized arithmetic mean of the available SGOV daily returns over a trailing 252-day window, 3.67%, and $\pi = 5\%$ is the premium assumed for one unit of market exposure. This is the equilibrium return of Sharpe (1964) and the starting point of Black and Litterman (1992), taken with no views of its own. Every stock is treated as fairly priced, and the solve spends all of its effort on risk.

The 5% premium is a chosen parameter. These are model-implied returns, not validated forecasts. Since SPY is both the market factor and an investable asset, its strong allocation is substantially explained by the model structure; it is not an independent empirical confirmation of CAPM.

2.4 Ten appetites

A portfolio is the long-only, fully invested set of weights w that maximizes

$$\mu^\top w - \frac{\gamma}{2} w^\top \Sigma w \quad \text{subject to} \quad w \geq 0, \mathbf{1}^\top w = 1, \quad (3)$$

for a risk aversion γ . We offer ten, $\gamma = 64, 32, 16, \dots, \frac{1}{8}$, so each successive point halves the penalty on variance. Risk and return need not double. Point 5 fixes $\gamma = 4$; its relationship to the tangency portfolio changes with the inputs.

3 An Active Set Solve

The usual way to solve (3) is to iterate. Take a step along the gradient, project back onto the weights that are non-negative and sum to one, and repeat. On a chain every step is paid for, and the steps crawl when some holding has almost no noise of its own, as an index fund does in a model built around the index. On this board, after 200,000 steps from an even start, the descent still held up to 41.7% of a portfolio in different places from the optimum, and its objective was lower at every one of the ten points. At 100 steps per point, a budget a contract could afford, up to 96.1% of a portfolio was misplaced. The one-factor structure allows something better.

Proposition 1. *Let D be diagonal with positive entries, $f \geq 0$ and $\gamma > 0$. Then (3) has a unique solution w^* . Writing $s = \beta^\top w$ for a portfolio's market exposure, every weight has the form*

$$w_i = \frac{\max(0, b_i - \lambda)}{\gamma d_i}, \quad b_i = \mu_i - \gamma f \beta_i s, \quad (4)$$

where, for each s , the level λ is the unique number that makes the weights sum to one. The function $h(s) = \beta^\top w(s) - s$ is strictly decreasing, with $h(\beta_{\min}) \geq 0 \geq h(\beta_{\max})$, and $w^* = w(s^*)$ at its root s^* .

Proof. Since D is positive definite, so is Σ , and (3) maximizes a strictly concave function over a convex set, so w^* is unique. Its Karush-Kuhn-Tucker conditions read $\mu - \gamma \Sigma w^* - \lambda \mathbf{1} + \nu = 0$ with $\nu \geq 0$ and $\nu_i w_i^* = 0$. By (1), $\Sigma w^* = D w^* + f \beta s^*$ with $s^* = \beta^\top w^*$, so a held stock satisfies $b_i - \lambda = \gamma d_i w_i^*$ and a stock left out satisfies $b_i - \lambda = -\nu_i \leq 0$, which is (4). For a fixed s , the weights (4) are the conditions of the separable problem: maximize $c^\top w - \frac{\gamma}{2} w^\top D w$ over the same set, with $c = \mu - \gamma f s \beta$. The sum $\sum_i \max(0, b_i - \lambda) / (\gamma d_i)$ falls strictly in λ wherever it is positive, so λ is unique. The solution of the separable problem is the gradient of a convex conjugate evaluated at c , a monotone map, so $(c' - c)^\top (w(c') - w(c)) \geq 0$. If $f = 0$, $w(s)$ is constant and $h(s)$ has slope -1 . For $f > 0$, raising s to s' moves c by $-\gamma f (s' - s) \beta$, hence $\beta^\top w(s') \leq \beta^\top w(s)$, and h falls strictly because of its $-s$ term. Last, $\beta^\top w$ is an average of the β_i with weights w , so it lies in $[\beta_{\min}, \beta_{\max}]$, which gives the signs of h at the ends. The root of h satisfies the conditions of (3), so it is w^* . $\square$

Once s is fixed, λ is found the way water finds its level across a row of tanks. Sort the b_i from high to low and admit stocks in that order while each one stands above the level that the stocks already admitted imply (Figure 2). One sort fixes λ , and s is found by bisection on h . The bisection only has to find which stocks are held. Once that set is fixed, the conditions $\mathbf{1}^\top w = 1$ and $\beta^\top w = s$ are linear in s and λ , and one more pass solves them exactly. Algorithm 1 is the whole solve for one appetite: twenty-five sorts of 192 numbers and a closing pass. Elton, Gruber and Padberg (1976) gave a ranking rule of the same family for the tangency portfolio of the single-index model; Proposition 1 carries the idea to every appetite on the dial.

The analytical closing step is exact in real arithmetic when the active set is correct. Twenty-four bisection steps do not guarantee active-set identification for every input. Fixed-point rounding adds approximation. The current closing pass operates on the candidate held set without a general KKT rejection-and-retry procedure, so the reported tests should not be read as a universal accuracy bound.




Figure 2. The water level for point 8 on 30 September 2026, where 28 stocks are held. Each bar is a stock’s b_i from equation (4), sorted from high to low. Stocks standing above the level are held, each in proportion to its height above the level divided by γd_i ; the rest are left out.

Algorithm 1 The morning solve for one appetite γ

- 1: $a_i \leftarrow 1/(\gamma d_i)$ for every stock; $lo \leftarrow \beta_{\min}$, $hi \leftarrow \beta_{\max}$
 - 2: **repeat** 24 times
 - 3: $s \leftarrow (lo + hi)/2$, $b_i \leftarrow \mu_i - \gamma f \beta_i s$
 - 4: sort the stocks by b , high to low
 - 5: admit stock k while $b_k \sum a > \sum ab - 1$, sums over the admitted stocks and k
 - 6: $\lambda \leftarrow (\sum ab - 1) / \sum a$, $w_i \leftarrow a_i \max(0, b_i - \lambda)$
 - 7: **if** $\beta^\top w > s$ **then** $lo \leftarrow s$ **else** $hi \leftarrow s$
 - 8: fill once more at $(lo + hi)/2$; the stocks with $w_i > 0$ are the held set S
 - 9: **return** w with s and λ solved exactly on S (Section 3.1)
-

3.1 Integer arithmetic

This implementation uses integer arithmetic throughout. Weights, exposures and the market’s variance carry 20 fraction bits; expected returns, own variances, the scores b_i and the level λ carry 40. The extra bits matter for a fund that is nearly all market, whose own variance is a few millionths: rounded to 20 bits it would lose most of its digits, and its weight would move with them. The closing pass writes the answer in a -weighted means and deviations,

$$s = \frac{C + \bar{\beta}}{1 + \gamma f V}, \quad \lambda = \bar{\mu} - \gamma f s \bar{\beta} - \frac{1}{\sum_S a_i}, \quad (5)$$

where $\bar{\beta}$ and $\bar{\mu}$ are the a -weighted means over the held set S , $C = \sum_S a_i(\beta_i - \bar{\beta})(\mu_i - \bar{\mu})$ and $V = \sum_S a_i(\beta_i - \bar{\beta})^2$. In this form it never subtracts two large sums.

The source file the contract compiles from is also compiled natively and tested. The following numerical results are reported for the tested inputs; they are not general bounds. On synthetic returns, at all ten points the held stocks earn the same marginal reward to within one part in a thousand of its scale, no stock left out earns more, and the weights sum to one within 0.00001. On the real board of 30 September 2026 the integer solve places every one of the ten portfolios within 0.01% of the double-precision answer, and the contract, run on a Nitro node, returns the same integers as the native build.

4 Cost

Table 1. Gas for a synthetic-data estimation and optimization benchmark over 192 assets on a local Nitro node running Stylus. The first row was sent as a transaction; the others are node estimates. The benchmark constructs synthetic inputs in memory and uses its own risk-aversion schedule, starting at 20,000. It does not replay the real-data portfolios, read live stock prices or execute vault trades.

Work	Gas	Of one transaction
Exact solve, 60 days of returns, ten points	1,226,229	3.8%
Exact solve, 120 days	1,308,535	4.1%
Exact solve, 252 days	1,497,223	4.7%
Projected gradient, 100 steps a point	6,467,536	20.2%
Full 192 by 192 covariance, Cholesky factor, two solves	9,058,852	28.3%

Robinhood Chain lets one transaction use up to 32,000,000 gas, a limit read from the chain’s `ArbGasInfo` precompile on 1 October 2026.¹ At that day’s gas price of 0.020056 gwei, the computation benchmark corresponds to 0.0000246 ETH, and a year of trading days 0.006 ETH for computation only. Only the estimation of β and d grows with the window; the solve does not, which is why a full year of history adds about a quarter of a million gas.

These figures cover the arithmetic. The sixty-day tape also has to be read from storage. Packed eight returns to a slot it is 1,440 slots, and at the standard 2,100 gas for a first read that is about 3.0 million gas more. Added to the benchmark, this is approximately 4.25 million gas, or 13.3% of the stated limit. This is an estimate, not an end-to-end measurement. Price publication, report verification, vault accounting and execution trades are additional. The gas price above is dated 1 October 2026.

5 The Board on 30 September 2026

Table 2. The ten portfolios solved from prices through 30 September 2026. Risk is the model’s yearly standard deviation; expected return follows equation (2). Point 5, with $\gamma = 4$, is shaded. Names counts weights above 0.1%; smaller positive positions may also exist.

Point	γ	Risk	Expected	Names	Largest holdings
1	64	0.7%	4.0%	4	SGOV 92%, SPY 6%
2	32	1.4%	4.3%	5	SGOV 86%, SPY 12%
3	16	2.8%	4.9%	5	SGOV 74%, SPY 25%
4	8	5.6%	6.2%	4	SPY 50%, SGOV 49%
5	4	11.2%	8.7%	2	SPY 99%, VTI 1%
6	2	22.0%	13.5%	135	SPY 45%, QQQ 5%
7	1	42.2%	22.3%	95	SOFI 4%, USAR 3%
8	1/2	57.3%	27.8%	28	USAR 8%, ALAB 8%
9	1/4	67.7%	30.1%	12	AEHR 16%, ALAB 15%
10	1/8	79.7%	31.7%	6	AEHR 29%, AAOI 17%

The calm end is the bill fund. Points 1 to 4 hold bills and SPY and nothing else of size: point 1 is 92% bills and 6% SPY, and point 4 is close to half and half. These allocations resemble the bills-and-market mix in Tobin’s framework. Point 5 holds 98.8% in SPY and 0.9% in VTI. SPY is the factor used to construct the model and its residual variance is floored as described in Section 2.2. This allocation is consistent with those assumptions. Point 5 is a fixed-risk-aversion portfolio, not a certified maximum-Sharpe portfolio.

At lower risk aversion a long-only investor cannot borrow to buy more of the market, so the bold-er points reach for stocks whose β is above one. Point 6 keeps 45% in SPY and spreads the rest across 134 other holdings. Point 10 holds 6 stocks with a combined risk of 80% a year (Figures 1 and 3).




Figure 3. What each point holds on 30 September 2026. Funds are the index, sector and commodity funds on the board. Points 1 to 4 mix bills with the index, point 5 is predominantly the index, and the bolder points move from funds into single companies.

6 A Year, Walked Forward

To see what the ten points would have done, we ran the solve on every trading day from 3 October 2025 to 30 September 2026, estimating weights with returns through close d and applying them to the close-to-close return from d to $d + 1$. This assumes fills at the same close used for estimation; it does not model the delay needed to observe that close and execute a trade. Prices are daily closes with dividends included, from Yahoo Finance’s public chart service.² Table 3 and Figure 4 show the result before trading costs.

Table 3. Walk-forward results over 248 trading days, before trading costs. Turnover proxy is half the sum of absolute changes between consecutive target weights, averaged per day, with an initial empty portfolio. It does not adjust prior holdings for price drift.

Portfolio	Return	Yearly risk	Worst fall	Turnover proxy
SPY	15.2%	13.1%	-8.9%	
Point 1	3.8%	0.7%	-0.3%	0.4%
Point 2	3.9%	1.4%	-0.8%	0.5%
Point 3	4.1%	2.9%	-1.8%	0.8%
Point 4	4.5%	5.7%	-3.9%	1.3%
Point 5	7.0%	11.3%	-8.3%	2.3%
Point 6	18.1%	22.1%	-12.5%	3.8%
Point 7	45.1%	43.5%	-22.1%	5.6%
Point 8	53.7%	67.8%	-42.6%	7.1%
Point 9	132.5%	80.5%	-49.0%	9.0%
Point 10	162.3%	91.7%	-49.7%	10.3%

SPY returned 15.2%, with a worst fall of 8.9%. Point 5 returned 7.0%. Its appetite is fixed, so when the market’s sixty-day variance rose it asked for less market and moved part of its money into bills, which also held its worst fall to 8.3%. Moreira and Muir (2017) studied this behaviour as the volatility-managed portfolio. The higher-numbered points bought high- β stocks in a year that paid for them: point 7 returned 45.1%, and point 10 returned 162.3% with a worst fall of 49.7% on the way. One year is one draw from the market, and the table reads as a description of that year.

The reported turnover proxy is 2.3% per day for point 5 and 10.3% for point 10. Multiplying it by 0.1% and 252 gives illustrative annual deductions of about 0.6% and 2.6%. These are not measured implementation costs. If 0.1% is charged on each dollar bought and each dollar sold, both sides must be counted, approximately doubling those estimates before accounting for portfolio drift and other execution effects.

The universe is selected as of 30 September 2026 rather than reconstructed at each historical date, so selection and survivorship effects remain possible. Eligibility requires a full sixty-return window. The simulation assigns zero to a missing next-day return and ignores target weights below one millionth. The reference calculations use floating-point bisection, with 30 iterations in the backtest and 40 for the snapshot, rather than replaying the integer contract on each date. Adjusted stock prices do not establish executable stock-token prices. A reproducible execution study needs archived inputs, delayed fills, drift-adjusted turnover, and sensitivity to spreads, slippage, and the variance floor.




Figure 4. \$1,000 held in four of the points and in SPY, 3 October 2025 to 30 September 2026, logarithmic scale. Weights use close d and are applied to the return from d to d+1; execution delay is not modeled. Before trading costs.

7 Proposed Vault Design

The proposed product assigns one vault to each point. A holder would deposit USDG and receive shares, ISO-1 to ISO-10, valued against the assets held by the vault. After each accepted solve, the vault would trade toward its new target weights. The design contemplates both USDG withdrawals and proportional withdrawals of the assets held. A proportional withdrawal can avoid liquidation trades, but still depends on the transferability and availability of the underlying tokens.

The intended design has no management fee and allocates eligible daily execution costs across vaults in proportion to their size. It proposes using Chainlink Data Streams reports to maintain a daily price tape. These are design choices; the computation benchmark does not verify their integration or demonstrate a deployed operating vault.

Before the design can be evaluated as a live system, its specification must identify execution venues, slippage limits, price-report verification and freshness rules, share-valuation timing, entry and withdrawal accounting, permissions, and recovery from partial or failed rebalances. Gas reimbursement also needs explicit eligibility and spending controls. This paper establishes neither an audit of these mechanisms nor evidence that every described vault function is currently deployed.

8 Conclusion

A single-factor covariance structure makes a family of mean-variance portfolios tractable in a constrained on-chain execution environment. The reported Stylus implementation computes ten portfo-

lios over a 192-token universe using 1,226,229 gas in a local estimation and optimization benchmark. An analytical closing step avoids relying solely on iterative convergence, while fixed-point tests provide evidence of numerical accuracy on the tested inputs. Persistent-storage reads and actual portfolio execution add costs beyond that benchmark.

The portfolio results reflect explicit assumptions: a chosen market premium, a sixty-day risk window, fixed risk-aversion coefficients, and long-only investment. The one-year simulation illustrates their behavior before trading costs. The next requirements for a reproducible and operational system are disclosure of numerical edge-case handling, reproducible data and code, end-to-end execution measurements, and validation of the proposed vault mechanisms. The contribution demonstrated here is computational feasibility under a stated model and benchmark, with live portfolio operation remaining a separate engineering claim.

Notes

¹ Address `0x...006C`, method `getGasAccountingParams`, which returned a limit of 32,000,000 gas per transaction and a speed limit of 7,000,000 gas per second.

² Daily bars from the version 8 chart endpoint, adjusted for dividends and splits. One of the 192 tokens, SATS, returned no history and is left out of every calculation.

References

- Black, F. and Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal* 48(5), 28 to 43.
- Boyd, S. and Johansson, K. (2024). Markowitz portfolio construction at seventy. arXiv:2401.05080.
- Elton, E. J., Gruber, M. J. and Padberg, M. W. (1976). Simple criteria for optimal portfolio selection. *Journal of Finance* 31(5), 1341 to 1357.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance* 7(1), 77 to 91.
- Michaud, R. O. (1989). The Markowitz optimization enigma: is “optimized” optimal? *Financial Analysts Journal* 45(1), 31 to 42.
- Moreira, A. and Muir, T. (2017). Volatility-managed portfolios. *Journal of Finance* 72(4), 1611 to 1644.
- Sharpe, W. F. (1963). A simplified model for portfolio analysis. *Management Science* 9(2), 277 to 293.
- Sharpe, W. F. (1964). Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance* 19(3), 425 to 442.
- Tobin, J. (1958). Liquidity preference as behavior towards risk. *Review of Economic Studies* 25(2), 65 to 86.

Appendix A Measured and Chosen

Table 4. Where each figure in this paper comes from.

Figure	Value	Source
Gas of the exact solve	1,226,229	Measured. Transaction on a local Nitro node, Stylus, status 1
Gas of the other methods	Table 1	Measured. Node estimates, same node and contract
Limit per transaction	32,000,000	Read from Robinhood Chain, 1 October 2026
Gas price	0.020056 gwei	Read from Robinhood Chain, 1 October 2026
Optimality of the integer solve	Section 3.1	Tested natively on the contract’s own source, synthetic and real boards; matching reported outputs on a local Nitro node
Ten points today	Table 2	Computed from daily closes through 30 September 2026
Year walked forward	Table 3	Computed from daily closes, before trading costs
Reading the tape from storage	about 3.0 million gas	Computed from the standard price of a first read
Window, premium, appetites	60 days, 5%, 64 to 1/8	Chosen
Vault terms and price source	Proposed vaults	Design choices; integration not demonstrated by the benchmark