

Demystifying Singular Covariance Matrices in Portfolio Optimization: A Rank-1 Inversion Approach for Financial Educators

September 26, 2026

Abstract

The classical Markowitz mean-variance optimization framework typically assumes a strictly positive definite, invertible covariance matrix. However, educators and researchers frequently encounter singular covariance matrices when dealing with high-dimensional asset sets over limited time horizons ($N > T$) or when assets exhibit perfect linear dependency. Standard textbook treatments often bypass this problem by deferring to abstract numerical algorithms or the Moore–Penrose pseudoinverse, which can obscure the underlying financial and geometric intuition. This paper provides a self-contained, analytically rigorous, and fully elementary guide to solving the Markowitz problem with a singular covariance matrix V of rank $N - 1$. We show that the null vector $\mathbf{z}$ spanning $\ker(V)$ can be extracted directly from any nonzero column of the classical adjugate matrix $\text{adj}(V)$, requiring nothing beyond cofactor expansion. Using $\mathbf{z}$, we introduce a rank-1 spectral shift $B = V + \mathbf{z}\mathbf{z}^T$ that renders the covariance matrix fully invertible while preserving its structure, yielding an explicit closed-form solution $\mathbf{w}^*(\mu_T)$ expressed entirely in terms of V , $\boldsymbol{\mu}$, $\mathbf{z}$, and the target return μ_T . We show that the resulting minimum-variance frontier collapses from a smooth hyperbola into a sharp, linear V-shape, where $\mathbf{z}$ acts as a synthetic zero-risk asset constructed endogenously from purely risky components.

1 Introduction

The mean-variance portfolio optimization framework established by Harry Markowitz [1] remains a cornerstone of quantitative finance education. In its standard formulation, computing the optimal allocation of assets requires inverting an empirical covariance matrix, denoted here by V . Standard treatments, including the classical two-fund separation results of Merton [2], assume that V is strictly positive definite ($V \succ 0$), guaranteeing that it is non-singular and invertible.

However, in both classroom environments and empirical research, this assumption frequently fails. The most common driver of this breakdown is the “curse of dimensionality”: when the number of assets N exceeds the number of available historical observation periods T , the sample covariance matrix lacks full rank and is structurally singular ($\det(V) = 0$). Alternatively, perfect linear dependencies among assets, such as a rigid currency peg or redundant derivative structures, will identically drop the determinant of V to zero. Bodnar and Schmid [5] and Pang [4] study aspects of this degeneracy formally, while Best and Grauer [3] examine the sensitivity of mean-variance solutions near such singularities.

When encountering $\det(V) = 0$, standard pedagogical approaches typically recommend three options:

1. Discarding assets until full rank is restored, which introduces selection bias and strips out valuable data.

2. Utilizing abstract optimization software or quadratic programming packages that solve the system as a numerical “black box.”
3. Deploying the Moore–Penrose pseudoinverse (V^+), which requires spectral machinery (eigen-decomposition or SVD) that many students have not yet mastered.

This note offers a direct algebraic alternative that uses only tools available in a first linear algebra course: determinants, cofactor expansion, and the classical adjugate matrix. A singular covariance matrix of corank one turns out to synthesize its own zero-risk asset out of purely risky components, and this asset can be extracted, and the full efficient frontier solved in closed form, using nothing more advanced than $\text{adj}(V)$.

2 Problem Setup and Notation

Consider a market with N risky assets. Let $\boldsymbol{\mu} \in \mathbb{R}^N$ denote the vector of expected asset returns, and let $V \in \mathbb{R}^{N \times N}$ represent the symmetric, positive semi-definite covariance matrix. The investor’s portfolio allocation is defined by the weight vector $\mathbf{w} \in \mathbb{R}^N$.

The classical mean-variance optimization problem seeks to minimize total portfolio variance subject to a target return constraint μ_T and a fully invested budget constraint:

$$\min_{\mathbf{w}} \quad \frac{1}{2} \mathbf{w}^T V \mathbf{w} \tag{1}$$

$$\text{s.t.} \quad \mathbf{w}^T \boldsymbol{\mu} = \mu_T \tag{2}$$

$$\mathbf{w}^T \mathbf{e} = 1 \tag{3}$$

where $\mathbf{e} \in \mathbb{R}^N$ is a vector of ones.

We assume throughout that V is singular with $\det(V) = 0$ and rank exactly $N - 1$ (a single zero eigenvalue; the case of higher corank is addressed in Section 10). By definition, $\ker(V)$ is one-dimensional, spanned by some nonzero vector.

3 Constructing the Null Vector via the Adjugate Matrix

Rather than obtaining the null vector via eigendecomposition, we extract it directly from the adjugate (classical adjoint) of V , using only cofactor expansion.

Lemma 1. *For any square matrix $V \in \mathbb{R}^{N \times N}$,*

$$V \cdot \text{adj}(V) = \text{adj}(V) \cdot V = \det(V) I. \tag{4}$$

Consequently, if $\det(V) = 0$, then every column of $\text{adj}(V)$ lies in $\ker(V)$.

Proof. This is the classical adjugate identity, following directly from the cofactor (Laplace) expansion of the determinant along any row or column. Setting $\det(V) = 0$ gives $V \cdot \text{adj}(V) = 0$, so each column of $\text{adj}(V)$ is annihilated by V . $\square$

Corollary 1. *If V is symmetric with rank exactly $N - 1$, then $\text{adj}(V)$ has rank exactly 1, and every nonzero column of $\text{adj}(V)$ is a scalar multiple of the (unique up to scale) null vector spanning $\ker(V)$.*

Proof. Since $\text{rank}(V) = N - 1$, at least one $(N - 1) \times (N - 1)$ minor of V is nonzero, so $\text{adj}(V) \neq 0$. By Lemma 1, every column of $\text{adj}(V)$ lies in $\ker(V)$, which is one-dimensional. A nonzero matrix all of whose columns lie in a one-dimensional subspace has rank exactly one, and every column is therefore a scalar multiple of the same vector. $\square$

Remark 1 (Diagnostic use). *Corollary 1 also gives a simple rank diagnostic: if every entry of $\text{adj}(V)$ is zero, then $\text{rank}(V) \leq N - 2$ ($\text{corank} \geq 2$), and the rank-1 construction of this paper does not directly apply; see Section 10.*

This yields a fully constructive algorithm for the null vector:

1. Compute $\text{adj}(V)$ via cofactor expansion (equivalently, $\text{adj}(V) = \det(V)V^{-1}$ in the limiting sense, though no inversion is actually performed since $\det(V) = 0$).
2. Select any nonzero column $\mathbf{c}$ of $\text{adj}(V)$.
3. Rescale so that the budget constraint holds: provided $\mathbf{e}^T \mathbf{c} \neq 0$, set

$$\mathbf{z} = \frac{\mathbf{c}}{\mathbf{e}^T \mathbf{c}}, \quad \text{so that } \mathbf{e}^T \mathbf{z} = 1. \quad (5)$$

Financially, $\mathbf{z}$ represents a fully-invested portfolio constructed entirely from risky assets that achieves absolute variance zero ($\mathbf{z}^T V \mathbf{z} = 0$). We denote its expected return by $\mu_z = \mathbf{z}^T \boldsymbol{\mu}$. The case $\mathbf{e}^T \mathbf{c} = 0$, in which no such rescaling exists, is discussed in Section 10.

4 The Method of Lagrange Multipliers

To find the optimal weight vector $\mathbf{w}^*$, we construct the Lagrangian with multipliers λ_1 and λ_2 :

$$L(\mathbf{w}, \lambda_1, \lambda_2) = \frac{1}{2} \mathbf{w}^T V \mathbf{w} - \lambda_1 (\mathbf{w}^T \boldsymbol{\mu} - \mu_T) - \lambda_2 (\mathbf{w}^T \mathbf{e} - 1). \quad (6)$$

Taking the gradient with respect to $\mathbf{w}$ and setting it to zero yields the first-order condition:

$$V \mathbf{w} = \lambda_1 \boldsymbol{\mu} + \lambda_2 \mathbf{e}. \quad (7)$$

Because V is symmetric, multiplying (7) from the left by $\mathbf{z}^T$ and using $\mathbf{z}^T V = \mathbf{0}^T$ gives:

$$\mathbf{z}^T V \mathbf{w} = \lambda_1 \mathbf{z}^T \boldsymbol{\mu} + \lambda_2 \mathbf{z}^T \mathbf{e} \quad (8)$$

$$0 = \lambda_1 \mu_z + \lambda_2. \quad (9)$$

Hence $\lambda_2 = -\lambda_1 \mu_z$, and substituting back into (7) collapses the system to a single multiplier:

$$V \mathbf{w} = \lambda_1 (\boldsymbol{\mu} - \mu_z \mathbf{e}). \quad (10)$$

5 The Rank-1 Regularized Inversion

Because $\det(V) = 0$, we cannot isolate $\mathbf{w}$ by multiplying by V^{-1} . We instead regularize V using the null vector $\mathbf{z}$ obtained in Section 3, defining

$$B = V + \mathbf{z} \mathbf{z}^T. \quad (11)$$

Lemma 2. $B = V + \mathbf{z}\mathbf{z}^T$ is strictly positive definite and invertible, with

$$B^{-1}\mathbf{z} = \frac{1}{\mathbf{z}^T\mathbf{z}}\mathbf{z}. \quad (12)$$

Proof. Since V is real symmetric with rank $N-1$, it admits an orthonormal eigenbasis $\{\hat{\mathbf{z}}, \mathbf{v}_2, \dots, \mathbf{v}_N\}$ with eigenvalues $\{0, \lambda_2, \dots, \lambda_N\}$, $\lambda_i > 0$, where $\hat{\mathbf{z}} = \mathbf{z}/\|\mathbf{z}\|$. Then $B\mathbf{z} = V\mathbf{z} + \mathbf{z}(\mathbf{z}^T\mathbf{z}) = (\mathbf{z}^T\mathbf{z})\mathbf{z}$, so $\mathbf{z}$ remains an eigenvector of B , with eigenvalue shifted from 0 to $\mathbf{z}^T\mathbf{z} > 0$. For $i \geq 2$, orthogonality gives $\mathbf{z}^T\mathbf{v}_i = 0$, so $B\mathbf{v}_i = \lambda_i\mathbf{v}_i$. All eigenvalues of B are strictly positive, so $B \succ 0$ and is invertible; (12) follows immediately from $B\mathbf{z} = (\mathbf{z}^T\mathbf{z})\mathbf{z}$. $\square$

Adding $\mathbf{z}\mathbf{z}^T\mathbf{w}$ to both sides of (10):

$$B\mathbf{w} = \lambda_1(\boldsymbol{\mu} - \mu_z\mathbf{e}) + (\mathbf{z}^T\mathbf{w})\mathbf{z}. \quad (13)$$

Multiplying by B^{-1} and applying (12):

$$\mathbf{w} = \lambda_1 B^{-1}(\boldsymbol{\mu} - \mu_z\mathbf{e}) + \left(\frac{\mathbf{z}^T\mathbf{w}}{\mathbf{z}^T\mathbf{z}}\right)\mathbf{z}. \quad (14)$$

Multiplying (14) from the left by $\mathbf{e}^T$ and imposing $\mathbf{e}^T\mathbf{w} = 1$, $\mathbf{e}^T\mathbf{z} = 1$:

$$\frac{\mathbf{z}^T\mathbf{w}}{\mathbf{z}^T\mathbf{z}} = 1 - \lambda_1 \mathbf{e}^T B^{-1}(\boldsymbol{\mu} - \mu_z\mathbf{e}). \quad (15)$$

Substituting back into (14):

$$\mathbf{w} = \mathbf{z} + \lambda_1 \underbrace{\left(B^{-1}(\boldsymbol{\mu} - \mu_z\mathbf{e}) - [\mathbf{e}^T B^{-1}(\boldsymbol{\mu} - \mu_z\mathbf{e})]\mathbf{z} \right)}_{=: \mathbf{h}}. \quad (16)$$

6 Closing the System via Target Return

Multiplying (16) from the left by $\boldsymbol{\mu}^T$ and using $\boldsymbol{\mu}^T\mathbf{w} = \mu_T$, $\boldsymbol{\mu}^T\mathbf{z} = \mu_z$:

$$\mu_T - \mu_z = \lambda_1(\boldsymbol{\mu} - \mu_z\mathbf{e})^T B^{-1}(\boldsymbol{\mu} - \mu_z\mathbf{e}) =: \lambda_1 D, \quad (17)$$

so

$$\lambda_1 = \frac{\mu_T - \mu_z}{D}, \quad D = (\boldsymbol{\mu} - \mu_z\mathbf{e})^T B^{-1}(\boldsymbol{\mu} - \mu_z\mathbf{e}). \quad (18)$$

Combining (18) with (16) delivers the closed-form optimal portfolio, expressed entirely in terms of V (through $B = V + \mathbf{z}\mathbf{z}^T$), $\boldsymbol{\mu}$, $\mathbf{z}$, and the target return μ_T :

$$\mathbf{w}^*(\mu_T) = \mathbf{z} + \left(\frac{\mu_T - \mu_z}{D}\right)\mathbf{h}, \quad \mathbf{h} = B^{-1}(\boldsymbol{\mu} - \mu_z\mathbf{e}) - [\mathbf{e}^T B^{-1}(\boldsymbol{\mu} - \mu_z\mathbf{e})]\mathbf{z} \quad (19)$$

7 Geometric Structure of the Risk Frontier

The achieved variance expands as

$$\text{Var}(\mathbf{w}^*) = (\mathbf{w}^*)^T V \mathbf{w}^* = (\mathbf{z} + \lambda_1 \mathbf{h})^T V (\mathbf{z} + \lambda_1 \mathbf{h}). \quad (20)$$

Since $\mathbf{z} \in \ker(V)$, all terms involving $V\mathbf{z}$ vanish:

$$\text{Var}(\mathbf{w}^*) = \lambda_1^2 (\mathbf{h}^T V \mathbf{h}). \quad (21)$$

Lemma 3. $V\mathbf{h} = \boldsymbol{\mu} - \mu_z \mathbf{e}$.

Proof. Write $\mathbf{u} = \boldsymbol{\mu} - \mu_z \mathbf{e}$. Note $\mathbf{z}^T \mathbf{u} = \mathbf{z}^T \boldsymbol{\mu} - \mu_z (\mathbf{z}^T \mathbf{e}) = \mu_z - \mu_z \cdot 1 = 0$. Since B^{-1} is symmetric, (12) gives $\mathbf{z}^T B^{-1} \mathbf{u} = (B^{-1} \mathbf{z})^T \mathbf{u} = \frac{1}{\mathbf{z}^T \mathbf{z}} \mathbf{z}^T \mathbf{u} = 0$. Then, using $V = B - \mathbf{z} \mathbf{z}^T$ and $\mathbf{h} = B^{-1} \mathbf{u} - (\mathbf{e}^T B^{-1} \mathbf{u}) \mathbf{z}$:

$$\begin{aligned} V\mathbf{h} &= VB^{-1} \mathbf{u} - (\mathbf{e}^T B^{-1} \mathbf{u}) V\mathbf{z} = VB^{-1} \mathbf{u} \\ &= (B - \mathbf{z} \mathbf{z}^T) B^{-1} \mathbf{u} = \mathbf{u} - \mathbf{z} (\mathbf{z}^T B^{-1} \mathbf{u}) = \mathbf{u} - \mathbf{z} \cdot 0 = \mathbf{u}. \end{aligned} \quad (22)$$

Using Lemma 3, $\mathbf{h}^T V \mathbf{h} = \mathbf{h}^T (\boldsymbol{\mu} - \mu_z \mathbf{e}) = D$ (matching (18)), so

$$\text{Var}(\mathbf{w}^*) = \left(\frac{\mu_T - \mu_z}{D} \right)^2 D = \frac{(\mu_T - \mu_z)^2}{D}, \quad \text{SD}(\mathbf{w}^*) = \frac{|\mu_T - \mu_z|}{\sqrt{D}}. \quad (23)$$

Rather than the standard hyperbolic Markowitz frontier, the frontier collapses into two straight rays meeting at the vertex $(0, \mu_z)$ in (SD, μ_T) -space, where portfolio risk drops to exactly zero. This gives an endogenous, purely risky-asset analogue of the Capital Market Line, without positing an external risk-free asset.

8 Worked Numerical Example

Let $N = 3$ with

$$V = \frac{1}{9} \begin{pmatrix} 8 & -10 & -20 \\ -10 & 50 & -50 \\ -20 & -50 & 200 \end{pmatrix}, \quad \boldsymbol{\mu} = \begin{pmatrix} -1/3 \\ 2/3 \\ 7/3 \end{pmatrix}, \quad \mathbf{e} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}. \quad (24)$$

One confirms $\det(V) = 0$ and $\text{rank}(V) = 2 = N - 1$. Direct cofactor expansion gives

$$\text{adj}(V) = \frac{1}{27} \begin{pmatrix} 2500 & 1000 & 500 \\ 1000 & 400 & 200 \\ 500 & 200 & 100 \end{pmatrix}, \quad (25)$$

whose three columns are, as guaranteed by Corollary 1, scalar multiples of one another. Taking the first column and rescaling by $\mathbf{e}^T \mathbf{c} = \frac{4000}{27}$:

$$\mathbf{z} = \begin{pmatrix} 5/8 \\ 1/4 \\ 1/8 \end{pmatrix}, \quad \mu_z = \mathbf{z}^T \boldsymbol{\mu} = \frac{1}{4}, \quad \mathbf{z}^T \mathbf{z} = \frac{15}{32}. \quad (26)$$

Applying (19) yields the closed-form efficient frontier

$$w_1^*(\mu_T) = \frac{10}{13} - \frac{15}{26} \mu_T, \quad w_2^*(\mu_T) = \frac{11}{65} + \frac{21}{65} \mu_T, \quad w_3^*(\mu_T) = \frac{4}{65} + \frac{33}{130} \mu_T, \quad (27)$$

and correspondingly $\text{SD}(\mathbf{w}^*(\mu_T)) = \sqrt{32/13} |\mu_T - \frac{1}{4}|$. At $\mu_T = \mu_z = 1/4$, (19) correctly recovers $\mathbf{w}^* = \mathbf{z} = (5/8, 1/4, 1/8)$, the zero-variance portfolio, as required.

9 Pedagogical Use

This material is best introduced after students have already met the standard, full-rank Markowitz problem and the classical two-fund separation result (typically following the treatment in Merton [2]), and before any discussion of shrinkage estimators or factor models. At that point, students already know the shape of the answer, a smooth hyperbola in mean-standard-deviation space, and the singular case gives a natural follow-up question: what happens to that shape when V cannot be inverted?

The derivation in Sections 3–8 was built to require nothing beyond a first course in linear algebra. Students who have computed a 3×3 or 4×4 adjugate by hand can reproduce every step, including the closed-form frontier, without any exposure to eigendecomposition. This also makes the result useful as a review exercise in an introductory linear algebra course itself: an instructor can use portfolio choice as a motivating application of the adjugate identity in Lemma 1, rather than presenting it as an isolated fact about determinants.

A natural in-class or homework exercise follows the structure of Section 8: give students a 3×3 (or, for a computational lab, 4×4) covariance matrix constructed to have rank $N - 1$ by design, for instance by defining one asset as an exact linear combination of the others, and a return vector μ . Students then

1. confirm $\det(V) = 0$ and $\text{rank}(V) = N - 1$;
2. compute $\text{adj}(V)$ by cofactor expansion, verify that its columns are proportional, and rescale to obtain $\mathbf{z}$;
3. form $B = V + \mathbf{z}\mathbf{z}^T$, invert it, and assemble $\mathbf{w}^*(\mu_T)$ from (19);
4. plot $\text{SD}(\mathbf{w}^*(\mu_T))$ against μ_T and confirm the frontier is two rays meeting at $(0, \mu_z)$, in contrast to the hyperbola obtained for a nearby full-rank perturbation of the same V .

The last step is particularly effective for building intuition: perturbing V slightly off the singular case and re-solving with the ordinary (full-rank) formula shows the hyperbola’s vertex approaching the kink point as the perturbation shrinks to zero, which ties the singular case directly to the standard result students already know rather than presenting it as a separate topic.

10 Limitations and Extensions

Two conditions are implicitly required for the method above to apply as stated:

(i) *Corank exactly one.* If $\text{rank}(V) \leq N - 2$, Corollary 1 shows $\text{adj}(V) \equiv 0$, so no null vector can be extracted this way; $\ker(V)$ is then more than one-dimensional and a single rank-1 shift $\mathbf{z}\mathbf{z}^T$ is insufficient to regularize V . The construction generalizes to a rank- k shift $B = V + \mathbf{Z}\mathbf{Z}^T$ using a basis $\mathbf{Z} \in \mathbb{R}^{N \times k}$ for $\ker(V)$, at the cost of the fully elementary adjugate route, since $\mathbf{Z}$ then requires either eigendecomposition or Gaussian elimination on V to isolate.

(ii) $\mathbf{e}^T \mathbf{c} \neq 0$. If every nonzero column $\mathbf{c}$ of $\text{adj}(V)$ satisfies $\mathbf{e}^T \mathbf{c} = 0$, no zero-variance, fully-invested portfolio exists, and the frontier does not collapse to the linear form of Section 6; the general (non-degenerate-frontier) singular case must instead be handled via the Moore–Penrose pseudoinverse or an equivalent bordered-Lagrangian system.

Both conditions are generic (they hold for almost all $V, \mathbf{e}$ satisfying $\text{rank}(V) = N - 1$) and are easily checked numerically before applying (19).

11 Conclusion

This note details an accessible, fully elementary framework for mean-variance portfolio optimization under a singular covariance matrix of corank one. By extracting the null vector directly from the classical adjugate matrix, and using it to construct an invertible rank-1 shift $B = V + \mathbf{z}\mathbf{z}^T$, the entire problem is resolved using only cofactor expansion and elementary matrix algebra, with no eigendecomposition, spectral theory, or numerical pseudoinverse routines required. Financial educators can use this route to build student intuition for synthetic zero-risk assets, the geometry of degenerate efficient frontiers, and the classical adjugate identity in a concrete, motivated setting.

References

- [1] H. Markowitz, “Portfolio Selection,” *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.
- [2] R. C. Merton, “An Analytic Derivation of the Efficient Portfolio Frontier,” *Journal of Financial and Quantitative Analysis*, vol. 7, no. 4, pp. 1851–1872, 1972.
- [3] M. J. Best and R. R. Grauer, “Capital Asset Pricing Compatible with Observed Market Value Weights,” *The Journal of Finance*, vol. 40, no. 1, pp. 85–103, 1985.
- [4] J. S. Pang, “A New and Efficient Algorithm for a Class of Portfolio Selection Problems,” *Journal of Financial and Quantitative Analysis*, vol. 15, no. 5, pp. 1105–1121, 1980.
- [5] T. Bodnar and W. Schmid, “Econometrical Analysis of the Sample Efficient Frontier,” *The European Journal of Finance*, vol. 15, no. 3, pp. 317–335, 2009.