

A Simple Expected-Utility Foundation for Mean-Variance Portfolio Choice

Karl Whelan

University College Dublin

May 2026

Abstract

Existing derivations of Markowitz's mean-variance portfolio from expected utility theory require quadratic utility, elliptically distributed returns, or approximately zero excess returns. This paper shows that none of these assumptions are necessary. Expanding utility around risk-free wealth eliminates a simultaneity problem in the first-order conditions and yields a simple quadratic optimization problem whose solution is proportional to the classical Markowitz portfolio. We also derive and empirically evaluate the third-order optimal portfolio, finding that its composition is essentially identical to the Markowitz portfolio in both a two-asset and a ten-industry dataset spanning a century of US returns.

1. Introduction

Mean-variance portfolio analysis remains one of the central tools of modern finance. Despite its widespread use, there has long been debate about how consistent it is with expected utility theory. Existing derivations of the Markowitz portfolio rely on one of three assumptions: quadratic utility, elliptically distributed returns, or approximately zero excess returns.

Quadratic utility can be used to establish the optimality of mean-variance portfolios (Ingersoll, 1987) but this implies increasing relative risk aversion and eventually declining utility as wealth rises. Similarly, the result can be proved under elliptically distributed returns, of which the normal distribution is the leading example (Chamberlain, 1983), but this rules out empirically important features such as skewness and fat tails. The Taylor-series approach rooted in the Arrow–Pratt literature avoids these assumptions, but standard derivations (such as Gollier, 2001) require assuming that expected excess returns are approximately zero.

This paper shows that none of these assumptions are necessary to derive Markowitz portfolios as approximately optimal under expected utility theory. Expanding utility around risk-free wealth yields a simple second-order optimization problem whose solution is proportional to the classical Markowitz portfolio. The derivation requires only smooth concave utility and finite return moments, and does not require any assumption about the distribution of returns or that mean excess returns are approximately zero. We also derive the third-order optimal portfolio and show empirically that its composition is essentially identical to the Markowitz portfolio, confirming that the second-order rule is robust to higher-order corrections.

2. Expected-Wealth Expansions and the Simultaneity Problem

Consider an investor with initial wealth w , risk-free gross return R_f , and a vector of scaled risky excess returns \mathbf{R}^e . End-of-period wealth from holding portfolio \mathbf{x} is

$$W = wR_f(1 + \mathbf{x}'\mathbf{R}^e) \quad (1)$$

Let $\boldsymbol{\mu}^e = \mathbb{E}[\mathbf{R}^e]$ and $\Sigma = \text{Cov}(\mathbf{R}^e)$.

The Arrow–Pratt literature on small risks suggests expanding expected utility as a Taylor series around expected wealth,

$$\mathbb{E}[W] = wR_f(1 + \mathbf{x}'\boldsymbol{\mu}^e) \quad (2)$$

A second-order expansion gives

$$\mathbb{E}[U(W)] \approx U(\mathbb{E}[W]) + \frac{1}{2}U''(\mathbb{E}[W])w^2R_f^2\mathbf{x}'\Sigma\mathbf{x} \quad (3)$$

where the first-order term vanishes by construction. At first glance this appears to deliver a mean-variance objective directly. However, both utility terms in (3) depend on portfolio choice through $\mathbb{E}[W]$. Differentiating fully with respect to \mathbf{x} gives

$$U'(\mathbb{E}[W])wR_f\boldsymbol{\mu}^e + U''(\mathbb{E}[W])w^2R_f^2\Sigma\mathbf{x} + \frac{1}{2}U'''(\mathbb{E}[W])wR_f\boldsymbol{\mu}^ew^2R_f^2\mathbf{x}'\Sigma\mathbf{x} = 0 \quad (4)$$

Even though the approximation truncates utility at second order, the first-order condition involves the third derivative of utility, because utility curvature depends on portfolio choice through $\mathbb{E}[W]$. Gollier (2001, Section 5.2) noted this simultaneity problem.

The standard resolution requires two steps, both invoking the assumption that $\boldsymbol{\mu}^e$ is small — not merely that its square or cube is small, but that mean excess returns are themselves close to zero. First, the third-order term in (4) is dropped on the grounds that it is of higher order under this assumption. Second — and separately — the utility derivatives $U'(\mathbb{E}[W])$ and $U''(\mathbb{E}[W])$ are evaluated at wR_f rather than $\mathbb{E}[W]$, which follows directly from $\boldsymbol{\mu}^e \approx 0$ making $\mathbb{E}[W] = wR_f(1 + \mathbf{x}'\boldsymbol{\mu}^e) \approx wR_f$.¹ Without this second step, the utility derivative coefficients still depend on \mathbf{x} and the system remains highly nonlinear even after dropping the third-order term. Applying both steps gives a linear first-order condition whose solution is

$$\mathbf{x}^* = -\frac{U'(wR_f)}{wR_fU''(wR_f)}\Sigma^{-1}\boldsymbol{\mu}^e = \frac{1}{\rho(w)}\Sigma^{-1}\boldsymbol{\mu}^e \quad (5)$$

where

$$\rho(w) = -\frac{wR_fU''(wR_f)}{U'(wR_f)} \quad (6)$$

is the coefficient of relative risk aversion evaluated at risk-free wealth. The approximation is valid as a local expansion in small mean excess returns, but it requires assuming that expected excess returns are approximately zero, an assumption that may not be empirically accurate.

¹Gollier (2001) uses a more formal version of this device. He introduces an auxiliary parameter k and writes the risky excess return as $\tilde{x} = k\mu + \tilde{y}$, where $\mathbb{E}[\tilde{y}] = 0$, then derives portfolio demand as a local approximation around $k = 0$. This avoids simply imposing $\boldsymbol{\mu}^e = 0$ by hand, but the economic content is the same: both steps of the argument are local to the zero-mean benchmark.

3. Expansion Around Risk-Free Wealth

Here we show that an alternative approach—expanding around the wealth the investor would have by holding only the risk-free asset—eliminates the simultaneity problem entirely and allows for an exact solution of the second-order problem without assuming approximately zero mean excess returns. We derive the optimal portfolio using this approach, explain how it still implies investors hold the Markowitz portfolio of risky assets and then discuss its implications.

3.1. The Main Result

A third-order Taylor expansion of $U(W)$ around wR_f gives

$$U(W) = U(wR_f) + U'(wR_f)wR_f(\mathbf{x}'\mathbf{R}^e) + \frac{1}{2}U''(wR_f)w^2R_f^2(\mathbf{x}'\mathbf{R}^e)^2 + \frac{1}{6}U'''(wR_f)w^3R_f^3(\mathbf{x}'\mathbf{R}^e)^3 \quad (7)$$

Since all utility terms are evaluated at the fixed point wR_f , taking expectations and dividing by the positive constant $U'(wR_f)wR_f$ gives the dimensionless objective

$$\mathbb{E}[\tilde{U}] = \mathbf{x}'\boldsymbol{\mu}^e - \frac{\rho(w)}{2}\mathbf{x}'M_2\mathbf{x} + \frac{\rho(w)\pi(w)}{6}\mathbb{E}[(\mathbf{x}'\mathbf{R}^e)^3] \quad (8)$$

where $M_2 = \mathbb{E}[\mathbf{R}^e\mathbf{R}^{e\prime}]$ is the matrix of raw second moments and

$$\pi(w) = -\frac{wR_fU'''(wR_f)}{U''(wR_f)} \quad (9)$$

is the coefficient of relative prudence (Kimball, 1990) evaluated at risk-free wealth. All preference parameters are constants independent of \mathbf{x} .

Retaining only the second-order terms gives the concave quadratic problem

$$\max_{\mathbf{x}} \quad \mathbf{x}'\boldsymbol{\mu}^e - \frac{\rho(w)}{2}\mathbf{x}'M_2\mathbf{x} \quad (10)$$

This has an exact first-order condition $\rho(w)M_2\mathbf{x} = \boldsymbol{\mu}^e$ and solution

$$\mathbf{x}_2 = \frac{1}{\rho(w)}M_2^{-1}\boldsymbol{\mu}^e \quad (11)$$

At first glance this appears to be a different portfolio from the classical Markowitz solution, since it involves the inverse of the raw second-moment matrix M_2 rather than the inverse of the covariance matrix Σ . However, the two matrices are related by the decomposition

$$M_2 = \Sigma + \boldsymbol{\mu}^e\boldsymbol{\mu}^{e\prime} \quad (12)$$

and even knowing this, it is not obvious that their inverses applied to $\boldsymbol{\mu}^e$ should produce portfolios with the same composition. The Sherman–Morrison identity, which applies precisely because $\boldsymbol{\mu}^e \boldsymbol{\mu}^{e'}$ is a rank-one matrix, reveals that they do:

$$M_2^{-1} \boldsymbol{\mu}^e = \frac{1}{1 + S_{\max}^2} \Sigma^{-1} \boldsymbol{\mu}^e \quad (13)$$

where S_{\max}^2 is the maximum squared Sharpe ratio achievable across all portfolios of the risky assets (Gibbons, Ross, and Shanken, 1989):

$$S_{\max}^2 = \boldsymbol{\mu}^{e'} \Sigma^{-1} \boldsymbol{\mu}^e \quad (14)$$

The M_2^{-1} portfolio is therefore not a different portfolio at all — it is exactly proportional to the classical Markowitz portfolio $\Sigma^{-1} \boldsymbol{\mu}^e$, scaled by the factor $1/(1 + S_{\max}^2)$. The optimal portfolio is

$$\mathbf{x}_2 = \frac{1}{\rho(w)(1 + S_{\max}^2)} \Sigma^{-1} \boldsymbol{\mu}^e \quad (15)$$

The derivation requires only smooth concave utility and finite return moments — no quadratic utility, no elliptical distributions, and no assumption that expected excess returns are small.

To get a sense of how large S_{\max}^2 might be in practice, we use two datasets, both spanning July 1926 to February 2026 (1,196 monthly observations) with the real risk-free rate constructed from Ken French’s one-month Treasury bill series (French, 2026) adjusted using CPI inflation from the Shiller (2026) dataset. The first uses real excess returns for US stocks and long-term government bonds from the Shiller (2026) dataset. The second uses real excess returns for the Ken French (2026) 10 Industry Portfolios. The two datasets give $S_{\max}^2 \approx 0.028$ and $S_{\max}^2 \approx 0.033$ respectively, confirming that the additional scaling factor $1/(1 + S_{\max}^2)$ is very close to one in both cases.

3.2. Implications

These results have several interesting implications.

First, this result delivers a version of Tobin’s (1958) two-fund separation theorem. The optimal portfolio allocates a fraction $1/\rho(w)(1 + S_{\max}^2)$ of wealth to the tangency portfolio $\Sigma^{-1} \boldsymbol{\mu}^e$ and the remainder to the risk-free asset. The expected-wealth expansion also produces the same separation, but only under the small-mean-excess-return assumption. Here that assumption is not required.

Second, the distinction between M_2 and Σ in the two solutions reflects a difference in what these two matrices measure. Σ measures risk as deviations from the portfolio’s own mean return — it is the natural measure for a relative problem: given that you are going to invest in risky assets, what is the best mix? M_2 measures risk as deviations from zero return — it is the natural measure for an absolute problem: how far should you depart from the safety of the risk-free asset? But the decision investors actually face is not just which mix of risky assets to hold, but precisely how much to invest

in risky assets relative to the alternative of staying in the risk-free asset. That is why M_2 emerges naturally from the expansion around risk-free wealth.

The Sherman–Morrison identity then delivers a clean reconciliation: the composition of the optimal risky portfolio is identical whether you solve the absolute or the relative problem. Σ and M_2 disagree on scale but agree on composition, and the scale disagreement is governed by exactly S_{\max}^2 — the squared Sharpe ratio of the opportunity set, which measures how attractive risky assets are relative to the safe alternative.

Third, substituting the optimal portfolio (15) into the second-order objective (10) gives the maximized expected utility

$$V^* = \frac{1}{2\rho(w)} \frac{S_{\max}^2}{1 + S_{\max}^2} \quad (16)$$

Maximized expected utility depends on the investment opportunity set solely through S_{\max} , the maximum achievable Sharpe ratio, scaled by the reciprocal of relative risk aversion. This provides a first-principles justification for the Sharpe ratio as the sufficient statistic for investor welfare: for any smooth concave utility function and without distributional assumptions, it is what investors are maximizing, up to the scaling factor $1/2\rho(w)$.

4. Higher-Order Corrections and Portfolio Robustness

A natural concern with the second-order approximation is whether omitted higher-order terms might materially alter portfolio choice. In the traditional expected-wealth expansion, controlling these terms requires assuming that mean excess returns are approximately zero. In the present framework, the situation is different. Because utility is expanded around the fixed benchmark wealth level wR_f , the higher-order terms take the form $\mathbb{E}[(\mathbf{x}'\mathbf{R}^e)^3]$, $\mathbb{E}[(\mathbf{x}'\mathbf{R}^e)^4]$, \dots and their accuracy depends on the higher moments of the optimal portfolio return being sufficiently moderate relative to its second moment, not on mean excess returns being small.

Rather than pursuing an analytical bound on these terms, we examine directly whether the third-order correction materially changes the optimal portfolio. Retaining the cubic term gives the third-order objective

$$Q_3(\mathbf{x}) = \mathbf{x}'\boldsymbol{\mu}^e - \frac{\rho(w)}{2}\mathbf{x}'M_2\mathbf{x} + \frac{\rho(w)\pi(w)}{6}\mathbb{E}[(\mathbf{x}'\mathbf{R}^e)^3] \quad (17)$$

Differentiating gives the first-order condition

$$\boldsymbol{\mu}^e - \rho(w)M_2\mathbf{x} + \frac{\rho(w)\pi(w)}{2}\mathbb{E}[(\mathbf{x}'\mathbf{R}^e)^2\mathbf{R}^e] = 0 \quad (18)$$

Writing the third-order solution as $\mathbf{x}_3 = \mathbf{x}_2 + \Delta$, a first-order approximation around \mathbf{x}_2 gives

$$\Delta \approx \frac{\pi(w)}{2} M_2^{-1} \mathbf{g}_2, \quad \mathbf{g}_2 = \mathbb{E} [(\mathbf{x}_2' \mathbf{R}^e)^2 \mathbf{R}^e] \quad (19)$$

If \mathbf{g}_2 is approximately aligned with $\boldsymbol{\mu}^e$, the third-order correction mainly rescales the Markowitz portfolio without changing its composition. More generally, third co-moments can tilt portfolio composition away from the second-order solution.

Moreover, because \mathbf{x}_2 maximizes the quadratic objective (10), its gradient with respect to \mathbf{x} is zero by construction. The welfare gain from moving from \mathbf{x}_2 to \mathbf{x}_3 is therefore proportional to $\|\Delta\|^2$ rather than $\|\Delta\|$ — an application of the envelope theorem. A portfolio already at the top of the quadratic objective loses very little by not adjusting for higher-order terms, even when those terms are individually non-negligible.

To check the size of the differences between second- and third-order optimal portfolios, we need to make some assumption about preferences. We assume $\pi(w) = \rho(w) + 1$, a property that holds exactly for all CRRA utility functions.

Table 1 reports results for both datasets. Industry shares are expressed as fractions of the total absolute second-order portfolio weight $\sum_j |x_{2j}|$, so differences between \mathbf{x}_2 and \mathbf{x}_3 columns reflect pure composition changes, with scale effects captured separately in the Scale column. In both datasets, the third-order correction primarily affects the scale of risky investment rather than the composition of the risky portfolio. In the two-asset case, cosine similarity between \mathbf{x}_2 and \mathbf{x}_3 exceeds 0.9996 across all levels of risk aversion, and the stock share of the risky portfolio changes by less than 2 percentage points. In the 10-industry case, cosine similarities remain above 0.987, with individual industry shares changing by at most 4 percentage points. Thus, even beyond the second-order approximation, the investor still holds approximately the Markowitz risky portfolio. Higher-order terms affect how much of that portfolio investors hold, but leave its composition essentially unchanged.

Table 1: Third-Order Portfolio Correction: Scale and Composition

ρ	Scale ratio	Cosine similarity	Stock share \mathbf{x}_2	Stock share \mathbf{x}_3
<i>Panel A: Two-asset portfolio (Shiller stock and bond data)</i>				
1	1.137	0.999645	40.22%	41.60%
2	1.103	0.999788	40.22%	41.29%
3	1.092	0.999829	40.22%	41.18%
5	1.082	0.999859	40.22%	41.09%
10	1.076	0.999880	40.22%	41.03%

<i>Panel B: Ten-industry portfolio (French 10 Industry Portfolios)</i>												
ρ	Scale	Cosine	NoDur	Durbl	Manuf	Enrgy	HiTec	Telcm	Shops	Hlth	Utils	Other
<i>Second-order shares (\mathbf{x}_2, all ρ)</i>												
All ρ	–	–	23.5	4.9	–9.9	9.8	5.6	6.6	2.6	11.8	2.3	–23.0
<i>Third-order shares (\mathbf{x}_3)</i>												
1	1.055	0.9877	25.0	6.8	–5.8	10.6	5.2	4.8	1.2	12.4	0.9	–27.3
2	1.041	0.9926	24.6	6.4	–6.8	10.4	5.3	5.3	1.5	12.2	1.3	–26.3
3	1.037	0.9940	24.5	6.2	–7.1	10.4	5.3	5.4	1.6	12.2	1.4	–25.9
5	1.033	0.9951	24.4	6.1	–7.4	10.3	5.4	5.5	1.7	12.2	1.5	–25.6
10	1.030	0.9959	24.3	6.0	–7.6	10.3	5.4	5.6	1.8	12.1	1.5	–25.4

Notes: Scale ratio is $\|\mathbf{x}_3\|_1/\|\mathbf{x}_2\|_1$. Cosine similarity is $(\mathbf{x}_2'\mathbf{x}_3)/(\|\mathbf{x}_2\|\|\mathbf{x}_3\|)$. Portfolio shares are computed as a percentage of total absolute weight. Sample: July 1926 to February 2026 (1,196 monthly observations).

5. Conclusion

This paper provides a simple expected-utility foundation for mean-variance portfolio choice that does not rely on quadratic utility, elliptical return distributions, or small-excess-return approximations. Expanding utility around risk-free wealth eliminates a simultaneity problem that arises when expanding around expected wealth and yields a simple fixed-coefficient quadratic optimization problem. The solution is proportional to the classical Markowitz portfolio, with the scaling for individual investors given by the reciprocal of the Arrow–Pratt coefficient of relative risk aversion evaluated at risk-free wealth. The third-order optimal portfolio is empirically almost identical in composition to the Markowitz portfolio across both a simple two-asset and a ten-industry asset universe, confirming the robustness of the second-order rule.

References

- [1] Chamberlain, Gary (1983). “A Characterization of the Distributions that Imply Mean-Variance Utility Functions,” *Journal of Economic Theory*, Volume 29, pages 185–201.
- [2] French, Kenneth R. (2026). Kenneth R. French Data Library. https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.
- [3] Gibbons, Michael R., Stephen A. Ross, and Jay Shanken (1989). “A Test of the Efficiency of a Given Portfolio,” *Econometrica*, Volume 57, pages 1121–1152.
- [4] Gollier, Christian (2001). *The Economics of Risk and Time*. MIT Press, Cambridge, MA.
- [5] Ingersoll, Jonathan E. (1987). *Theory of Financial Decision Making*. Rowman & Littlefield, Totowa, NJ.
- [6] Kimball, Miles S. (1990). “Precautionary Saving in the Small and in the Large,” *Econometrica*, Volume 58, pages 53–73.
- [7] Sherman, Jack and Winifred J. Morrison (1950). “Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix,” *Annals of Mathematical Statistics*, Volume 21, pages 124–127.
- [8] Shiller, Robert J. (2026). Online Data: US Stock Markets 1871–Present and CAPE Ratio. <http://www.econ.yale.edu/~shiller/data.htm>.
- [9] Tobin, James (1958). “Liquidity Preference as Behaviour Towards Risk,” *Review of Economic Studies*, Volume 25, pages 65–86.