

Portfolio optimization: the **efficient frontier** in four formulas

Given what you expect each asset to return and how the assets move together, which mixes give the most return for each level of risk? XL Edge answers with live formulas: a covariance estimate, the two portfolios that define the frontier, the frontier itself with its chart, and the risk and return of any mix.

Functions `fx.RiskCovEstimate` · `fx.RiskEfficientPortfolios` · `fx.RiskFrontier` · `fx.RiskPortfolio`

In XL Edge XL Edge tab > Productivity Tools > Stat Tools > Insert Portfolio

Example *XL Edge Example - Portfolio.xlsx* (same folder). Every number in this guide is read from it.

Method Markowitz mean-variance, Merton's closed-form frontier; short sales allowed (no constraints).

1 The idea, and its one big catch

2 Using the ribbon

3 Inputs: expected returns and covariance

4 The covariance estimators: CovEstimate

5 GMVP and tangency: EfficientPortfolios

6 The frontier and its chart: Frontier

7 Any mix: Portfolio

8 Practical advice

9 Argument reference

1 The idea, and its one big catch

A portfolio's expected return is the weighted average of its assets' expected returns, but its risk (standard deviation) is *less* than the weighted average whenever the assets do not move in lockstep. That gap is diversification, and mean-variance optimization exploits it. For every target return there is one mix with the lowest possible risk; plotted, those mixes trace the **minimum-variance frontier**, and its upper half is the **efficient frontier**. Two portfolios on it matter most:

- the **global minimum-variance portfolio (GMVP)**, the least risky mix of all, at the frontier's left tip; and

- the **tangency portfolio**, the mix with the highest Sharpe ratio (excess return over the risk-free rate per unit of risk). Combining it with cash traces the **capital market line**, which beats every other portfolio on the frontier.

THE CATCH: GARBAGE IN, OPTIMIZED GARBAGE OUT

The optimizer believes its inputs exactly. Small errors in expected returns, which are notoriously hard to estimate, swing the weights wildly and favor whichever assets happened to look best in your sample. The example below shows it plainly. Treat the output as a disciplined starting point, use the covariance estimators that tame noise, and think hard about the expected returns (section 8).

2 Using the ribbon

XL Edge tab > Productivity Tools > Stat Tools (arrow) > Insert Portfolio

Menu item	XL Edge asks for	Result
Covariance Estimate	The returns table (one asset per column, one period per row, names in the header row above); the estimator; the market returns (single-index) or the shrinkage weight.	$(k + 2) \times (k + 1)$ labeled matrix
Efficient Portfolios (GMVP, Tangency)	The expected returns (one per asset), the $k \times k$ covariance matrix (numbers only), the risk-free rate (or <code>none</code>).	$4 \times (k + 4)$
Efficient Frontier (+ chart)	As above, plus the number of points along the frontier.	$(\text{Points} + 2) \times 7$, chart beside it
Portfolio Mean and SD	The weights, the expected returns, the covariance matrix.	2×3

Each writes one live formula after telling you its size. Change a return in the data and every result, the chart included, updates. The usual chain is: returns table → *Covariance Estimate* → *Efficient Frontier*, pointing at the numbers inside the covariance block.

3 Inputs: expected returns and covariance

The example has ten years of monthly returns for five asset classes and a market index (simulated, with typical but invented means, volatilities and correlations). Expected returns are simply the historical means, a common but *weak* choice, used here to show its weakness:

```
one mean per column of returns (a row of 5)
=BYCOL(Returns!B5:F124, LAMBDA(c, AVERAGE(c)))
```

Mean monthly return, 2016 - 2025 (simulated), sheet 1 *Covariance*

US equity	Intl equity	Bonds	Real estate	Commodities
0.92%	1.27%	0.43%	1.17%	0.04%

The data were generated with US equity expected to beat international equity, and commodities to earn about 5% a year. Ten years of noise reversed the first and nearly erased the second. The optimizer cannot know that.

4 The covariance estimators: `fx.RiskCovEstimateA`

Four ways to estimate the covariance matrix (Benninga, *Financial Modeling*, ch. 10). All keep the sample variances on the diagonal; they differ in the covariances, which are where the noise is:

Method	Covariances	US with Intl	When
"Sample" (default)	The sample covariances ($n - 1$).	0.002289	Long histories, few assets.
"SingleIndex"	$\beta_i \beta_j \times \text{var}(\text{market})$, each β the asset's slope on the market.	0.002511	Many stocks driven by one market factor.
"ConstantCorr"	One average correlation $\times \text{SD}_i \times \text{SD}_j$ (Elton and Gruber).	0.000814	A robust, simple baseline.
"Shrink"	Shrinkage $\times$ sample + $(1 - \text{Shrinkage}) \times$ its diagonal; 0.5 by default.	0.001144	Pull noisy covariances toward zero.

```
=fx.RiskCovEstimateA>Returns!B5:F124, , , , Returns!B4:F4)
=fx.RiskCovEstimateA>Returns!B5:F124, "SingleIndex", Returns!G5:G124, , Returns!B4:F4)
```

Look at the US / international covariance. The two equity markets really are highly correlated (0.75 in the simulation). The constant-correlation estimate forces them down to the average correlation of *all* pairs, which here includes bonds and commodities, a reminder that a "robust" estimator is a modeling choice, not a free lunch. The single-index estimate, driven by a market index built mostly from the two, keeps the link. Compare the weights each produces before you trust any of them.

The block is $(k + 2) \times (k + 1)$: a title, a header of names, then the matrix with names down the side. Portfolio functions want the *numbers only*, the $k \times k$ inside, here `'1 Covariance'!B12:F16`.

5 GMVP and tangency: `fx.RiskEfficientPortfoliosA`

```
expected returns, covariance (numbers only), risk-free rate 0.25% a month, names
=fx.RiskEfficientPortfoliosA('1 Covariance'!B6#, '1 Covariance'!B12:F16, B5, Returns!B4:F4)
```

Monthly figures; sheet 2 *Efficient*

Portfolio	Mean	SD	Sharpe	US eq.	Intl eq.	Bonds	Real est.	Commod.
GMVP	0.386%	1.487%	0.092	8.1%	-5.3%	89.1%	-0.3%	8.5%
Tangency	1.139%	3.800%	0.234	-32.9%	73.0%	62.1%	22.6%	-24.8%

Reading it. The GMVP is mostly bonds (89.1%), as the least volatile asset, with small positions elsewhere that diversify it further; its risk is 1.487% a month (5.15% a year, against 4.63% expected return). The tangency portfolio expects 13.67% a year with 13.16% risk, but look at how it gets there: **73.0% international equity and a 32.9% short in US equity**, plus a short in commodities. That is the catch from section 1 in action. Two highly correlated assets whose sample means differ by chance invite the optimizer to buy one and short the other.

NEGATIVE WEIGHTS ARE SHORT SALES

These functions solve the unconstrained problem: weights sum to 100% but may be negative. That is the textbook frontier and exact in closed form. If you cannot short, use the results as a guide, or find the long-only optimum with Excel's Solver (minimize `fx.RiskPortfolioλ`'s variance subject to weights ≥ 0 and a target mean).


The Sharpe ratio is $(\text{Mean} - R_f) / \text{SD}$. The tangency portfolio exists on the efficient frontier only when the risk-free rate is below the GMVP's mean. With no risk-free rate, only the GMVP row is filled.

6 The frontier and its chart: `fx.RiskFrontierλ`

Everything a frontier chart needs, in one spill: 50 points (by default) along the whole minimum-variance frontier, each with its SD and mean, and beside them the label, SD and mean of every asset, the GMVP, the tangency portfolio and the two ends of the capital market line. *Efficient Frontier (+ chart)* writes it and draws the chart.

```
=fx.RiskFrontierλ('1 Covariance'!B6#, '1 Covariance'!B12:F16, '2 Efficient'!B5, , Returns!B4:F4)
```

Efficient frontier (risk-free rate 0.25%)



XL Edge's Efficient Frontier chart, sheet *3 Frontier*. The orange curve is the minimum-variance frontier; its upper half is efficient. The dashed capital market line runs from the risk-free rate through the tangency portfolio (black diamond).

Reading it. Every asset sits to the right of the curve: some mix always gives the same return for less risk. Points on the lower half are dominated; you would never hold them. Investors who can borrow and lend at the risk-free rate do best on the capital market line, choosing their point by how much risk they

want. The frontier comes from Merton's closed form ($SD^2 = (C m^2 - 2 A m + B) / D$), so it is exact and recalculates instantly.

7 Any mix: `fx.RiskPortfolioλ`

Mean, variance and standard deviation of any set of weights: $w^T \mu$ and $w^T \Sigma w$.

```
=fx.RiskPortfolioλ(B6:F6, '1 Covariance'!B6#, '1 Covariance'!B12:F16)
```

Equal weights (20% each) give a mean of 0.764% and an SD of 3.414% a month. Feeding in the GMVP's weights returns its mean and SD exactly; the difference in the check cell on sheet *4 Mix* is 0.0e+00. Use it to price any proposal: your current allocation, a policy portfolio, a rounded version of the optimizer's weights.

8 Practical advice

- **Expected returns drive everything.** Historical means over 10 years are mostly noise. Prefer forward-looking estimates: equilibrium (CAPM-implied) returns, consensus forecasts, or historical means shrunk toward a common value.
- **Tame the covariance.** With many assets or short histories, use `"Shrink"` or `"SingleIndex"`; compare the weights across estimators, and distrust any that flip.
- **Keep units consistent.** Monthly returns give monthly means, SDs and risk-free rate. Annualize at the end: mean $\times 12$, SD $\times \sqrt{12}$ (sheet *2 Efficient* shows both).
- **Check the weights before the chart.** Large shorts, or weights far above 100%, signal input noise rather than insight.
- **Rf below the GMVP mean,** or there is no sensible tangency portfolio.
- **Pair with simulation.** Bootstrap the returns (see the bootstrap guide) and recompute the weights to see how unstable they are.

9 Argument reference

fx.RiskCovEstimate(Returns, [Method], [Market], [Shrinkage], [Names], [Labels])

Argument	Default	Meaning
Returns	required	n × k returns, one asset per column.
Method	"Sample"	"Sample", "SingleIndex", "ConstantCorr" or "Shrink".
Market	-	n × 1 market returns, for "SingleIndex".
Shrinkage	0.5	Weight on the sample matrix for "Shrink", 0 to 1.
Names, Labels	-, TRUE	Asset names; Labels = FALSE returns the bare k × k matrix.

fx.RiskEfficientPortfolios(Means, Cov, [Rf], [Names])

Argument	Default	Meaning
Means	required	Expected returns, one per asset (a row or a column).
Cov	required	k × k covariance matrix, numbers only.
Rf	none	Risk-free rate, in the same units. Without it, no tangency row and no Sharpe ratios.
Output, 4 × (k + 4): title; Portfolio / Mean / SD / Sharpe / one column per asset; GMVP; Tangency.		

fx.RiskFrontier(Means, Cov, [Rf], [Points], [Names])

Argument	Default	Meaning
Points	50	Points along the frontier (at least k + 4).
Output, (Points + 2) × 7: title; Point / SD / Mean along the frontier; Label / SD / Mean for the assets, GMVP, tangency and the capital market line's ends (#N/A without Rf, which charts skip).		

fx.RiskPortfolio(Weights, Means, Cov, [Labels]) **and** **fx.RiskEnvelope**(Means, Cov, C)

Argument	Default	Meaning
Weights	required	One weight per asset. Output: Mean, Variance, SD (2 × 3 with a header).
C	required	Envelope portfolio (Benninga's Proposition 1): weights proportional to $\text{Cov}^{-1}(\text{Means} - C)$. $C = R_f$ gives the tangency portfolio.

References: Markowitz (1952), *Journal of Finance* 7(1); Merton (1972), *JFQA* 7(4); Sharpe (1963), *Management Science* 9(2); Elton and Gruber (1973), *Journal of Finance* 28(5); Benninga, *Financial Modeling*, 4th ed., ch. 9-10. The example data are simulated for illustration. XL Edge and the fx.Risk... functions: © 2026 W Phillips, edgewisedata.com, MIT License.