60 Years of Portfolio Optimisation (2014)

- Markowitz (1952) was the first to quantify the risk/return tradeoff as an optimisation problem
- Classical framework must be modified for practical use

Standard MVO

- Universe of assets $S_1, S_2, \dots S_n$ with future returns $\mathbf{r} = [r_1, \dots r_n]^T$
- $\Omega \in \mathbb{R}^n$ is the set of all possible portfolios
- A portfolio is repr. by weights $\boldsymbol{\omega} = [\omega_1, \dots, \omega_n]^T$, $\boldsymbol{\omega} \in \Omega$
- Risk is encapsulated in the cov matrix
  - ↳ $\sigma_{ij} = \sigma_{ji} = \rho_{ij} \sigma_i \sigma_j$ for $i \neq j$
  - ↳ all cov matrices are positive semidefinite
  - ↳ if assets are linearly independent, Σ will be positive definite so Σ is invertible
  - ↳ portfolio volatility is given by $\sigma(\omega) = \sqrt{\omega^T \Sigma \omega}$

variances

$$\Sigma = \begin{bmatrix} \sigma_{11} & \cdots & \sigma_{1n} \\ \vdots & \ddots & \vdots \\ \sigma_{n1} & \cdots & \sigma_{nn} \end{bmatrix}$$

- $\mathbf{r}$ may be estimated by its mean $\boldsymbol{\mu}$, where $\mu_i = E(r_i)$
- The MVO problem can then be formulated:

$$\max_{\omega \in \Omega} \mu^T \omega - \lambda \cdot \omega^T \Sigma \omega$$

$$\max_{\omega \in \Omega} \mu^T \omega \quad \text{s.t.} \quad \omega^T \Sigma \omega \leq \sigma^2_{max}$$

$$\max_{\omega \in \Omega} \omega^T \Sigma \omega \quad \text{s.t.} \quad \mu^T \omega \geq R_{min}$$

- If the constraint set only includes linear equalities/inequalities, MVO is a quadratic program (QP)

Enhancing MVO

Transaction costs

- It is important to consider tx costs in allocation otherwise optimised portfolios may lose too much to slippage/commission.
- Slippage is the difference between the anticipated price at t_0 and the VWAP over $[t_0, t_0 + T]$, due to random fluctuations and market impact
- Slippage is most heavily influenced by trade size and asset liquidity
  ↳ one model is the Almgren model, which considers the permanent and temporary impacts for each order of x_i shares of stock i.

$$I^{perm}(x_i) = \gamma \cdot T \cdot \sigma_i \cdot \operatorname{sign}(x_i) \cdot \left| \frac{x_i}{V_i \cdot T} \right|^{\alpha} \cdot \left(\frac{\theta_i}{V_i} \right)$$

float

daily volume

$$I^{temp}(x_i) = \eta \cdot \sigma_i \cdot \operatorname{sign}(x_i) \cdot \left| \frac{x_i}{V_i \cdot T} \right|^{\beta}$$

fraction of the day over which trade executes

↳ looking at a large dataset of trades gives $\underbrace{\alpha \simeq 1 \quad \beta \simeq \tfrac{1}{2}}$

Almgren actually uses $\beta \simeq 3/5$

- Tx cost = amount traded × slippage:

$$TC_i(x_i) = |x_i| \cdot \left(\tfrac{1}{2} I^{perm}(x_i) + I^{temp}(x_i) \right)$$
$$= a_i x_i^2 + b_i \cdot |x_i|^{3/2}$$

- Let h_0 denote the dollar values of holdings, h denote new holdings, x be the vector of trades (# shares), p be price vector.

- The optimisation problem then becomes:

$$\max_{w \in \Omega} \quad \mu^T h - \lambda \cdot h^T \Sigma h - \gamma \cdot TC(x) \qquad TC(x) = \sum_i TC_i(x_i)$$

$$\text{s.t.} \quad h - x \circ p = h_0 \quad \leftarrow \text{new holdings come from trades}$$

$$h^T \mathbb{1} + TC(x) \leq h_0^T \mathbb{1} \quad \leftarrow \text{self-financing, i.e. mkt impact funded by selling holdings.}$$

- This is now a nonlinear program, either requiring special SOCP solvers or a QP relaxation.

Constraints in portfolio construction

- MVO allows users to incorporate many constraints, which may actually improved ex-post performance.
- Regulatory requirements must be respected even if they reduce performance - e.g short selling/leverage restrictions.
- Discretionary exposure constraints limit exposure to certain risk factors - these act as model insurance, reducing the effects of estimation errors.
- Trading constraints e.g "don't trade $> x\%$ of daily volume" may reduce tx costs.

Quantifying the effects of constraints

- The transfer coefficient is defined as the correlation coeff. between the risk-adjusted active weights in an optimised portfolio and the forecasted alphas of portfolio securities.

- Can be used to evaluate the effects of constraints.
  - ↳ for unconstrained portfolios with uncorrelated alphas, weights ∝ alphas.
  - ↳ but it does not decompose the effects of individual constraints.
- Shadow cost decomposition (using Lagrange multipliers) can ascribe an opportunity loss to each constraint
  - ↳ can be extended to attribute returns to objective terms.
  - ↳ evaluation is most useful on an ex-post basis.

Misalignment from constraints

- When the alpha model contains factors not present in the risk model, the optimiser may be unstable and underestimate true risk.
- Even when aligned, in the presence of weight bounds, risk may be greatly underestimated. An alpha alignment factor may be required to improve ex-post performance.

Improving Estimation Errors

- Classical MVO ignores estimation errors, and has 'error maximisier' properties.
- As a rule of thumb: expected return $\overset{2\times\text{ impt.}}{>}$ variance $\overset{2\times\text{ impt.}}{>}$ covariance
- A simple method is to use weight constraints, but this can impact stability.

- Alternatively, we can use diversification indicators as constraints
  ↳ related to information content: high information ⇒ concentrated.

Black-Litterman (BL)

- Expected rets are a weighted avg of market equilibrium and investor views, with weights depending on asset volatility/covariance and confidence.
  ↳ can also allow for relative views
- BL assumes that asset returns are multivariate normally-distributed, i.e $r \sim N(\mu, \Sigma)$ but μ itself is distributed as $\mu \sim N(\pi, \Sigma_\pi)$ where π is a vector of eq. returns.

1. Investor views are expressed as $P\mu \sim N(q, \Omega)$:
   $P \in \mathbb{R}^{k \times n}$ 'picks out' assets you have a view on
   $q \in \mathbb{R}^k$ expected return on views (alphas)
   $\Omega \in \mathbb{R}^{k \times k}$ cov matrix of views (confidence)
2. Market equilibrium is based on the CAPM:
   $$\pi_i = E(r_i) - r_f = \beta_i (E(r_m) - r_f)$$
   $$\beta_i = \frac{Cov(r_i, r_m)}{\sigma_m^2}$$
   - benchmark (mcap) weights: $w_b = [w_{b1}, \ldots, w_{bn}]^T$

- then the CAPM can be expressed as

$$\pi = \delta \Sigma w_b \quad \longleftarrow \text{vector of risks, including covariances}$$

$$\delta = \frac{E(r_m) - r_f}{\sigma_m^2} \quad \Big\} \text{ market price of risk}$$

3. BL expected returns

$$\hat{\mu}_{BL} = [(\tau\Sigma)^{-1} + P^T\Omega^{-1}P]^{-1} \underbrace{[(\tau\Sigma)^{-1}\pi + P^T\Omega^{-1}q]}_{\text{weighted avg of mkt with views}}$$

some tuning hyperparam

↳ if no views ($q=0$) or zero conf ($\Omega = 0$), $\hat{\mu}_{BL} = \pi$

↳ $(\tau\Sigma)^{-1}$ and $P^T\Omega^{-1}P$ are our conf in mkt and view.

- BL can be applied directly to any normally distributed prior (or zero for active managers)
- Extensions of BL can be used to incorporate views on volatility, nonlinear/non-normal views of returns etc.

Robust Optimisation

- Models optimisation problems with uncertainty sets on the parameters, then maximising the worst-case utility.
  - ↳ can work for other risk paradigms like VaR/CVaR.
  - ↳ should lead to satisfactory portfolios on most realisations of the parameters.
- Computationally tractable via second-order cone opt.
- It is unclear whether robust optimisation outperforms shrinkage.

Higher moments and tail risk

- Financial return distributions tend to have fat tails and asymmetries which cannot be described by mean-variances.
- It is possible to maximise utility under the empirical return dist., but generally MVO is a good approx (except for S-shaped utility functions).
- Skew and Kurtosis can be directly incorporated, as a better approx for general utility maximisation.
- CVaR optimisation can be formulated as a LP.

New directions in portfolio optimisation

Diversification

- 'Naive' $1/N$ diversification outperforms many policies out of sample - it is not subject to estimation error or data mining.
- Alternatively, we can optimise with only a risk model to find the global min variance (GMV) portfolio.

Measuring risk contributions

- One approach is to define the risk contrib. of position i as: $\hat{\sigma}_i(\omega) = \sigma(\omega) - \sigma(\omega_{-i})$ where ω_{-i} is the portfolio with the ith weight set to zero.

- However, this is unintuitive because $\sum_i \hat{\sigma}_i(w) \neq \sigma(w)$
- Alternatively, we define the marginal risk contrib. (MRC) of asset i as: $MRC_i(w) = \dfrac{\partial \sigma(w)}{\partial w_i}$ ← rate of change of risk as weight of i increases

$$(\sigma(w))^2 = w^T \Sigma w \implies 2\sigma(w) \nabla_w \sigma(w) = 2\Sigma w$$

$$\therefore MRC_i(w) = \frac{(\Sigma w)_i}{\sigma(w)} \leftarrow \text{ith component}$$

↳ the risk contrib. (RC) is then $RC_i(w) = w_i \cdot MRC_i(w)$

↳ note that $\sum_i RC_i(w) = \sigma(w)$

↳ the relative risk contrib (RRC) is $RRC_i(w) = \dfrac{RC_i(w)}{\sigma(w)}$

Risk Parity

- A portfolio is a risk parity portfolio with respect to Σ iff $RRC_i(w) = 1/n$, $i = 1, \dots, n$.
  - ↳ i.e total risk is allocated evenly across assets.
- In general, existence / uniqueness / constructions may be difficult or impossible depending on constraints.
- As an opt. problem, we aim to minimise the deviation from risk parity (DRP). Multiple measures could be used, e.g:

$$DRP(w) = \sum_i \sum_j \left[w_i (\Sigma w)_i - w_j (\Sigma w)_j \right]^2 \leftarrow \text{all pairwise diffs}$$

$$DRP(w) = \sum_i \left(\frac{w_i (\Sigma w)_i}{[\sigma(w)]^2} - \frac{1}{n} \right)^2 \leftarrow \text{sum of square deviations}$$

↳ but these metrics are nonconvex in w, so may get stuck in local minima and are generally harder to optimise.

- In the long-only case, we can solve:

$$\min_{\omega \geq 0} \; \omega^T \Sigma \omega - \sum_i \ln \omega_i \longrightarrow \text{logarithmic barrier}$$

↳ the optimality condition is $2\Sigma\omega - \omega^{-1} = 0$ (set grad = 0)

$\omega^{-1} \equiv \left[\frac{1}{\omega_1}, \frac{1}{\omega_2}, \cdots \frac{1}{\omega_n}\right]^T$

or equivalently $\omega_i (\Sigma\omega)_i = \frac{1}{2}$

↳ then $RRC_i(\omega) = {}^{1/2}/_{n/2} = \frac{1}{n}$

↳ since the optimality condition for this objective and risk parity are the same, we can just use this objective, which is convex in ω.

↳ results may need to be scaled such that $\sum_i \omega_i = 1$.

Mixing different sources of alpha

- Rather than combining multiple alpha views into a single value, it may be desirable to let model mixing occur at the optimiser level.
  - ↳ may want to constrain risk contribs due to each alpha source
  - ↳ alphas may have different periods!
- Consider two sets of alphas, e.g strategic (μ^X) and tactical (μ^Y), with respective active weights ω_A^X and ω_A^Y. Optimise:

$$\max_{\omega \in \Omega} \; (\mu^X)^T \omega_A^X + (\mu^Y)^T \omega_A^Y - \lambda \cdot (\omega_A^X + \omega_A^Y)^T \Sigma (\omega_A^X + \omega_A^Y) - \gamma TC(\omega)$$

$$\text{s.t.} \quad \omega - \omega_A^X - \omega_A^Y = \omega_B \qquad \text{weight = active + benchmark}$$

$$\left.\begin{aligned} (\omega_A^X)^T \Sigma \omega_A^X &\leq u^X \\ (\omega_A^Y)^T \Sigma \omega_A^Y &\leq u^Y \end{aligned}\right\} \text{budget risk between the two.}$$

↳ budgets can change based on performance

Views on groups of securities

- How should you allocate if you are bullish on AAPL but bearish on tech stocks?
- Similar to model mixing except we must map between groups and securities. This can be done with a $m \times n$ incidence matrix G, where $G_{gi} = 1$ if asset i is in group g, and zero otherwise.
  - ↳ benchmark security weights → group weights via $\omega_B^G = G\omega_B^I$
- Suppose we have group weights ω^G. We then scale the security weights in the group by $\omega_g^G / (\omega_B)_g^G$, i.e our group weight vs benchmark group weight:
  - ↳ this produces a security-level benchmark
  - ↳ we can then define the group/security active weights

$$\omega_A^G = \omega^G - \omega_B^G$$

$$\omega_A^I = \omega^I - \hat{\omega}_B^I \quad \leftarrow \text{defined w.r.t } \underline{\text{scaled}} \text{ benchmark}$$

- We can then optimise as before.

Multi-period Optimisation (MPO)

- MPO jointly models risk, alpha and its decay, and impact costs.
  - ↳ i.e *when* you should trade, not just *what*
- Let returns be modelled by $r_{t+1} = \mu_t + \alpha_t + \varepsilon_{t+1}^r$

returns to comp. for risk → (μ_t)

predictable excess alpha ↑ (α_t)

- Alphas are forecasted with a factor model with k mean-reverting factors:

$$\alpha_t = B f_t + \varepsilon_t^{\alpha}$$

$$\Delta f_{t+1} = -D f_t + \varepsilon_{t+1}^{f}$$

$\leftarrow B \in \mathbb{R}^{n \times k}$ factor loadings
$\leftarrow f \in \mathbb{R}^{k}$ factors

idiosyncratic components.

$D \in \mathbb{R}^{k \times k}$ mean reversion coefficients

- We can incorporate permanent and temporary tx costs by adding additional costs to the investor's α.
- The MPO problem is then given by:

diagonal tx cost matrix

$$\max_{\Delta w_1, \Delta w_2, \ldots \Delta w_{T-1}} E\left[\sum_{t=1}^{T-1} (1-\rho)^t \left(w_t^T \alpha_t - \frac{\lambda}{2} w_t^T \Sigma w_t - \frac{1}{2} \Delta w_t^T \Lambda w_t \right) + (1-\rho)^T \left(w_T^T \alpha_T - \frac{\lambda}{2} w_T^T \Sigma w_T \right) \right]$$

discount factor

↳ i.e. maximise PV of period returns less tx costs over every possible rebalance

- This is a stochastic linear-quadratic regulator problem, and can be solved with standard theory.