Lorenzo Naranjo October 30, 2026

The Treynor-Black Model

Most investors cannot analyze every security. A research department covers a few dozen stocks at most, and for the thousands of others the best guess is that the market prices them fairly. Treynor and Black (1973) showed how such an investor should use the analysts' work. Hold the market portfolio, which needs no forecasts, and add an **active portfolio** of the securities the analysts believe are mispriced, sized according to how good the forecasts are. The model rests on the single-index model of The Index Model, which keeps the inputs manageable: an alpha, a beta and a residual variance for each analyzed security, and the expected return and variance of the market.

Model Setup

We have N risky assets with excess returns following a single-index model

$$R_i = \alpha_i + \beta_i R_M + e_i,$$

where $\mathrm{E}(e_i) = 0$, $\mathrm{Cov}(R_M, e_i) = 0$ and $\mathrm{Cov}(e_i, e_j) = 0$ for $i \neq j$, for all $i, j = 1, \ldots, N$. Suppose we form a portfolio P with the N risky assets and the market portfolio M.

Denote by w_i the weights of the N risky assets and w_M the weight in the market portfolio. Of course, the sum of the portfolio weights must be equal to one, i.e.,

$$\sum_{i=1}^{N} w_i + w_M = 1. \tag{1}$$

We can write the excess returns of this portfolio over the risk-free asset as

$$R_P = \alpha_P + \beta_P R_M + e_P.$$

The portfolio alpha includes the alpha of all the securities except the market, which has no alpha, and the beta of the portfolio is a weighted average of the betas of all securities and the beta of the market, which is one:

$$\alpha_P = \sum_{i=1}^{N} w_i \alpha_i, \qquad \beta_P = \sum_{i=1}^{N} w_i \beta_i + w_M. \tag{2}$$

We will see later that the alpha of the optimal portfolio is positive, since its weights are positive for securities with positive alpha and negative otherwise. Choosing w_M given the weights of the other assets is the same as choosing β_P, so we will choose the weights in the individual assets and the beta of the portfolio.

The idiosyncratic risk includes the firm-specific risks of all securities except the market, which has none. Because the firm-specific risks are uncorrelated with each other, we have that

$$\sigma^2(e_P) = \sum_{i=1}^{N} w_i^2 \sigma^2(e_i). \tag{3}$$

Thus,

$$\mathrm{E}(R_P) = \sum_{i=1}^{N} w_i \alpha_i + \beta_P \, \mathrm{E}(R_M), \qquad \sigma_P^2 = \beta_P^2 \sigma_M^2 + \sum_{i=1}^{N} w_i^2 \sigma^2(e_i). \tag{4}$$

The expected return and variance in (4) depend on the alphas generated by the active securities and the systematic exposure of the resulting portfolio. This is the essence of the Treynor-Black model. The active part of the portfolio generates an alpha but adds diversifiable risk. It may also change the portfolio's beta, and with it the systematic risk and the expected return that comes with it.

Solving the Model

The objective is to find the portfolio weights $w_1, w_2, \ldots, w_N$ and w_M that maximize its Sharpe ratio, which is equivalent to finding $w_1, w_2, \ldots, w_N$ and the beta of the portfolio.

Maximizing the Sharpe ratio is equivalent to minimizing the variance of the portfolio for a given expected return. Only one particular expected return satisfies the restriction in (1). For other targets of the expected return, we would have to borrow or invest the

difference at the risk-free rate. If we denote by μ_P a possible target expected return, the minimization problem is

$$\min_{\{w_1,w_2,\ldots,w_N,\beta_P\}} \frac{1}{2}\sigma_P^2 \quad \text{s.t.} \quad \mathrm{E}(R_P) = \mu_P - r_f.$$

With a Lagrange multiplier λ, the first-order conditions for the weights and the beta are

$$w_i\sigma^2(e_i) - \lambda\alpha_i = 0, \quad i = 1, \ldots, N, \qquad \beta_P\sigma_M^2 - \lambda\,\mathrm{E}(R_M) = 0.$$

The condition on the expected return is not important, since we will compute λ later by imposing that the weights add up to one. The first condition implies that for each security

$$w_i = \lambda\frac{\alpha_i}{\sigma^2(e_i)}. \tag{5}$$

The weight of each security increases with its alpha but decreases with the variance of its firm-specific risk. In other words, the weight of each security depends on the *quality* of the alpha. A small alpha with little firm-specific variance can be a better opportunity than a large alpha with a large firm-specific variance. Assets with negative alphas carry a negative weight, whereas assets with positive alphas have a positive weight.

Using the weights of the individual assets, the weight in the market portfolio is

$$w_M = 1 - \lambda\sum_{i=1}^{N}\frac{\alpha_i}{\sigma^2(e_i)}. \tag{6}$$

The second condition implies that

$$\beta_P = \lambda\frac{\mathrm{E}(R_M)}{\sigma_M^2}. \tag{7}$$

On the other hand, the definition of the portfolio beta in (2) together with (6) gives

$$\beta_P = \lambda\sum_{i=1}^{N}\frac{\alpha_i}{\sigma^2(e_i)}\beta_i + w_M = \lambda\sum_{i=1}^{N}\frac{\alpha_i}{\sigma^2(e_i)}(\beta_i - 1) + 1. \tag{8}$$

Combining (7) with (8) we get

$$\lambda = \frac{1}{\frac{\mathrm{E}(R_M)}{\sigma_M^2} + \sum_{i=1}^{N} \frac{\alpha_i}{\sigma^2(e_i)}(1-\beta_i)}. \tag{9}$$

The previous equation provides an expression for λ that can be computed using the model inputs. The value for λ can then be used to compute the weights of the individual assets and of the market portfolio.

The Active Portfolio

Equation (5) says that the securities enter in fixed proportions to each other, $\alpha_i/\sigma^2(e_i)$, whatever the value of λ. The solution can therefore be read in two steps.[1] First, collect the analyzed securities into a single active portfolio A with weights

$$w_i^A = \frac{\alpha_i/\sigma^2(e_i)}{\sum_{j=1}^{N} \alpha_j/\sigma^2(e_j)}, \tag{10}$$

which add up to one. The active portfolio is itself a security with an alpha, a beta and a residual variance:

$$\alpha_A = \sum_{i=1}^{N} w_i^A \alpha_i, \qquad \beta_A = \sum_{i=1}^{N} w_i^A \beta_i, \qquad \sigma^2(e_A) = \sum_{i=1}^{N} \left(w_i^A\right)^2 \sigma^2(e_i). \tag{11}$$

Second, combine A with the market. Applying (5) and (9) with a single security, the active portfolio, gives

$$w_A^0 = \frac{\alpha_A/\sigma^2(e_A)}{\mathrm{E}(R_M)/\sigma_M^2}, \qquad w_A^* = \frac{w_A^0}{1+(1-\beta_A)w_A^0}, \tag{12}$$

with $1 - w_A^*$ in the market. The first expression is the position in A if its beta were one: it compares the alpha of A per unit of residual variance with the market's risk premium per unit of variance. The second adjusts for the beta of A. Part of a position in A is market exposure, which replaces part of the position in the market. When $\beta_A > 1$, the investor

[1] Treynor and Black (1973) derive the optimal portfolio differently and describe it as three stages: select an active portfolio, blend it with the market to maximize the Sharpe ratio, and scale the combination up or down by lending or borrowing. Only the last stage depends on the investor's attitude toward risk. The two steps here are their first two stages, read off the single optimization above.

Table 1: Forecasts for three stocks, their weights in the active portfolio A and in the optimal portfolio P.

	Alpha (%)	Beta	Residual vol. (%)	Weight in A (%)	Weight in P (%)
Ashford	3.0	1.2	30.0	49.0	16.1
Brightwater	1.5	0.8	15.0	98.0	32.3
Calder	−2.0	1.0	25.0	−47.1	−15.5
Market		1.0			67.1

holds more of A and correspondingly less of the market, and when $\beta_A < 1$, less of A. Substituting (11) into (12) gives back the weights $w_i = w_A^* w_i^A$ of (5) and (9).

Example 1 (An Active Portfolio of Three Stocks). The market has an expected excess return of 8% and a volatility of 20%. A research team covers three stocks, with the forecasts in Table 1. Ashford has the largest alpha, but also the largest residual risk. Brightwater has half the alpha with half the residual volatility, so a quarter of the residual variance, and gets the larger weight. Calder is overpriced and is sold short.

By (11), the active portfolio has an alpha of 3.88%, a beta of 0.90 and a residual volatility of 23.9%. By (12), $w_A^0 = (0.0388/0.0571)/(0.08/0.04) = 0.340$ and $w_A^* = 0.340/(1 + (1 - 0.902) \times 0.340) = 0.329$. The investor puts 32.9% of the portfolio in A and 67.1% in the market. The optimal portfolio has a beta of 0.97, an expected excess return of 9.02% and a volatility of 20.89%, for a Sharpe ratio of 0.432 against the market's 0.400. □

Figure 1 shows the geometry. The curve traces the portfolios that combine A and the market in different proportions. The active portfolio alone has a low Sharpe ratio, because most of its risk is residual, but mixed with the market its residual risk is diluted while its alpha still adds to the expected return. The optimal portfolio P is the point of the curve with the steepest capital allocation line, which lies above the market's.

The Information Ratio

We can use (5) and (7) in (4) to get the expected return and the variance of the optimal portfolio:

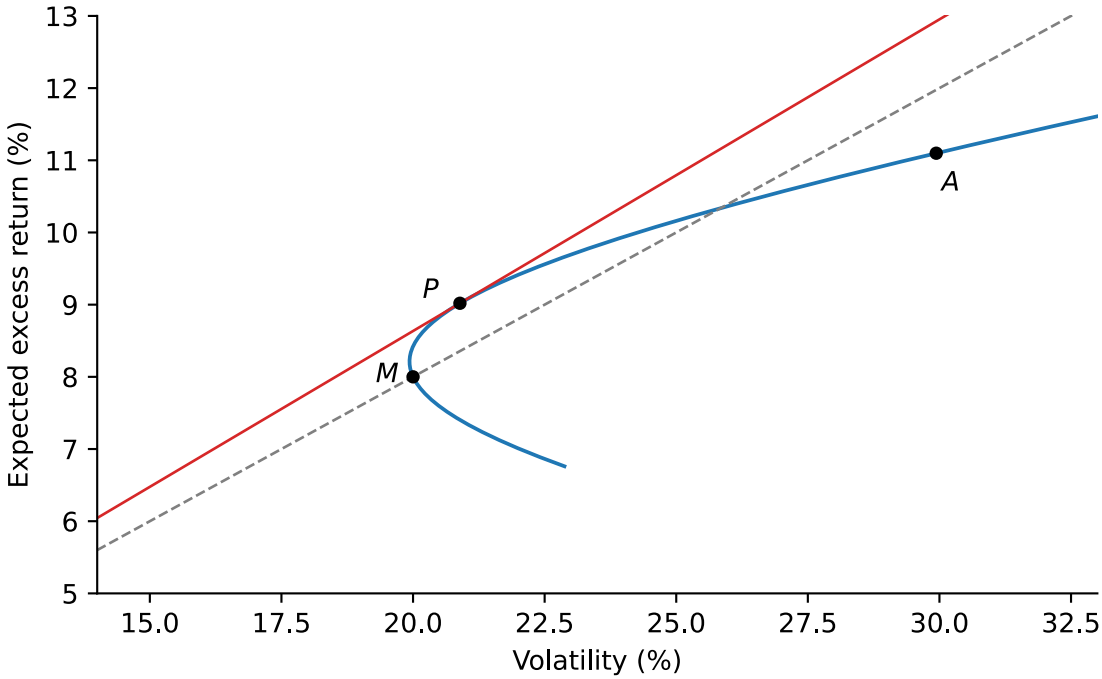


Figure 1: Expected excess return and volatility of portfolios that combine the active portfolio A of Example 1 with the market M. The optimal portfolio P has the steepest capital allocation line (solid). The dashed line is the capital allocation line of the market.

$$\mathrm{E}(R_P) = \lambda K, \qquad \sigma_P^2 = \lambda^2 K, \qquad K = \sum_{i=1}^{N} \left(\frac{\alpha_i}{\sigma(e_i)} \right)^2 + \left(\frac{\mathrm{E}(R_M)}{\sigma_M} \right)^2. \tag{13}$$

The Sharpe ratio is therefore $\sqrt{K}$, which gives the following relationship between the Sharpe ratio of the optimal portfolio and the market Sharpe ratio,

$$\left(\frac{\mathrm{E}(R_P)}{\sigma_P} \right)^2 = \sum_{i=1}^{N} \left(\frac{\alpha_i}{\sigma(e_i)} \right)^2 + \left(\frac{\mathrm{E}(R_M)}{\sigma_M} \right)^2. \tag{14}$$

The equation shows that the Sharpe ratio of the optimal portfolio is higher than the Sharpe ratio of the market alone. Each security contributes positively, regardless of whether its alpha is positive or negative. Securities with negative alphas have negative weights, i.e., they are sold short in order to take advantage of the fact that they are expensive relative to the model.

For a given security i, the term $\alpha_i/\sigma(e_i)$ is its **information ratio**, introduced in The Index Model, and it measures how *good* the alpha is. Treynor and Black call the alpha the appraisal premium and this ratio the **appraisal ratio**. If the volatility of the firm-specific risk is small, then it will be easier to diversify that risk away and hence extract the alpha of the security. Each security contributes to increasing the squared value of the Sharpe ratio of the optimal portfolio by the square of its information ratio. The active portfolio of (10) collects all of these contributions: its own squared information ratio, $\alpha_A^2/\sigma^2(e_A)$, is the sum of the squared information ratios of its securities. In Example 1, the information ratios are 0.10, 0.10 and −0.08, and (14) gives a Sharpe ratio of $\sqrt{0.10^2 + 0.10^2 + (-0.08)^2 + 0.4^2} = 0.432$, as before.

Since the contributions add up, many weak ideas can be worth as much as a few strong ones. With N independent forecasts that each have an information ratio of IR, the Sharpe ratio is $\sqrt{S_M^2 + N\,IR^2}$, where S_M is the Sharpe ratio of the market. Figure 2 plots it. A single stock with an information ratio of 0.1 barely moves the Sharpe ratio of the market, but 25 of them raise it from 0.40 to 0.64. The information ratio of each forecast matters as much as their number: the same 25 forecasts with an information ratio of 0.05 raise the Sharpe ratio only to 0.47.

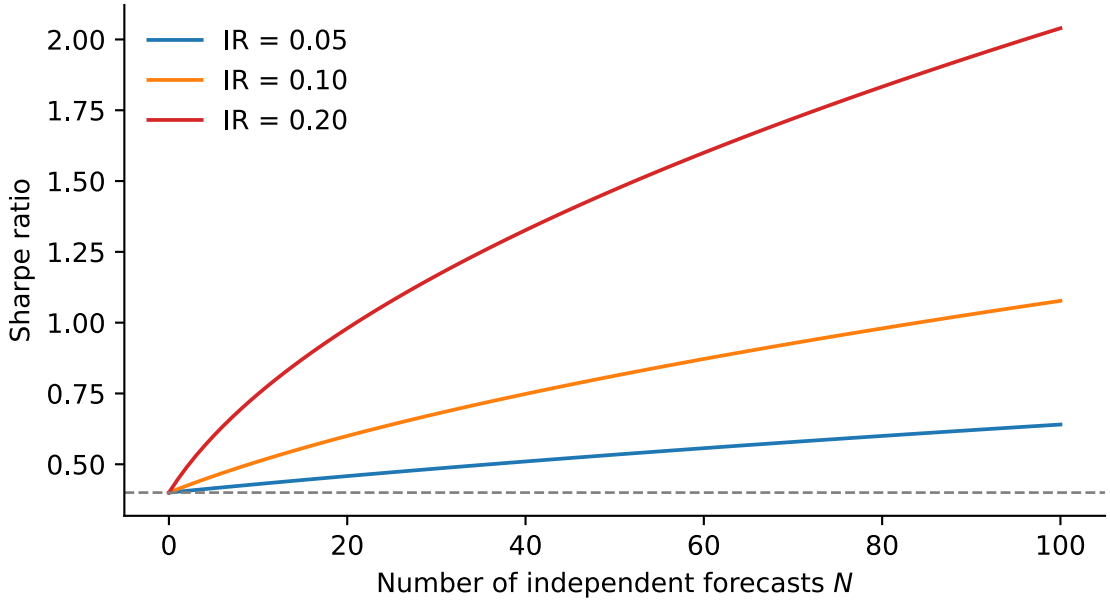


Figure 2: Sharpe ratio of the optimal portfolio with N independent forecasts that each have the same information ratio, when the market has a Sharpe ratio of 0.40.

Adjusting Alpha Forecasts

The model takes the alphas as known, but they are forecasts, and forecasts are noisy. Plugging a raw forecast into (5) treats it as if it were the true alpha. Suppose an analyst's forecast is the true alpha plus an independent error, $\alpha_i^f = \alpha_i + \varepsilon_i$, and that across stocks the true alphas average zero. The best estimate of the true alpha given the forecast is then

$$\hat{\alpha}_i = \frac{\sigma_\alpha^2}{\sigma_\alpha^2 + \sigma_\varepsilon^2}\,\alpha_i^f, \tag{15}$$

which shrinks the forecast toward zero, more so the noisier the forecasts are relative to the true dispersion of alphas. The shrinkage factor is the slope of a regression of true alphas on forecasts, and it can be estimated from the analyst's record: regress the alphas that were realized on the forecasts that were made, and use the slope. Realized alphas are noisy too, but their noise is unrelated to the forecasts and does not bias the slope. Treynor and Black (1973) propose exactly this regression of realized returns, net of the market, on the analyst's past estimates. Its intercept corrects an analyst who is persistently too optimistic or too pessimistic, and its slope one whose estimates are too extreme or too timid. This is the same logic as Estimation Risk and Shrinkage, applied to the alphas of the Treynor-Black model.

The adjustment matters more than it may seem. A forecast that is shrunk by a factor of 0.2 gets a fifth of the weight, and its contribution to the squared Sharpe ratio, the square of its information ratio, falls by a factor of 25. A raw forecast that is mostly noise does worse than nothing: as Problem 6 shows, it adds residual risk that no alpha compensates.

Historical Alphas as Forecasts

A tempting shortcut is to skip the analysts and use each stock's alpha from an index-model regression on past returns, as in The Index Model. Each January from 1932 to 2025, we take the 100 largest U.S. stocks with returns in each of the previous 60 months and estimate their alphas, betas and residual variances over those months. The forecast

is the past alpha. The realized alpha is the stock's average return in excess of $\beta_i R_M$ over the following year.

Figure 3 sorts the 9,385 stock-years into 20 groups by their forecast and plots the average realized alpha of each group. The average forecasts of the groups range from −14% to 35% a year, but the realized alphas show no relation to them, if anything a slightly negative one. The slope of the regression of realized on forecast alphas is −0.10, and the correlation −0.05. The slope is negative in 7 of the 10 decades. A large past alpha is mostly luck, and the stocks that did best over five years tend to do slightly worse afterwards.

The consequence for the Treynor-Black portfolio is severe. From 1972 to 2025, we form the optimal portfolio each January from the raw forecasts, with the market's expected excess return estimated by its average since 1926 and its variance over the previous 60 months, and hold it for the year. Past alphas look like excellent forecasts in the sample

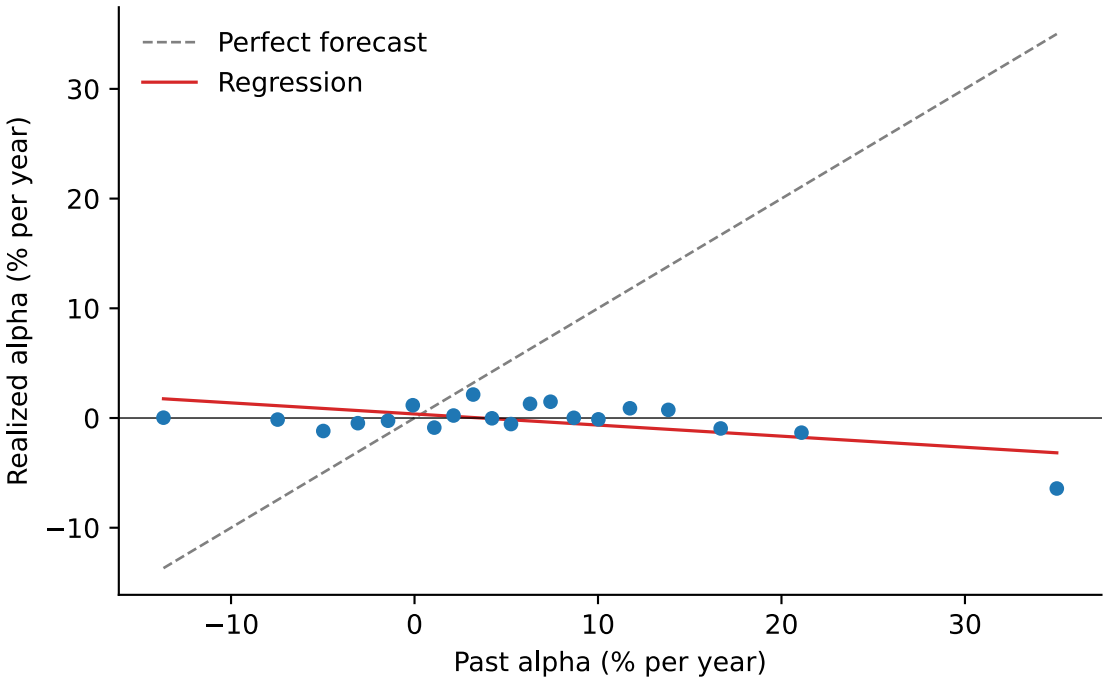


Figure 3: Average realized alpha in the following year against the past 60-month alpha, for the 100 largest U.S. stocks each January from 1932 to 2025, in 20 groups sorted by past alpha. The dashed line is a perfect forecast and the solid line the regression of realized on past alphas across all stock-years. Annualized.

Table 2: Annualized excess returns of the market and of Treynor-Black portfolios of the 100 largest U.S. stocks formed each January from past 60-month alphas, 1972 to 2025.

	Mean (%)	Volatility (%)	Sharpe ratio	Worst month (%)
Market	7.57	15.77	0.48	−23.19
Raw forecasts	42.57	172.28	0.25	−385.19
Adjusted forecasts	7.57	15.77	0.48	−23.19

that produced them, with large information ratios, so the model takes enormous active positions. Table 2 summarizes the results. In the median year, the active positions add up to 7.5 times the portfolio's value in absolute terms, financed in part by a short position in the market: the median weight in the market is −3.7. The portfolio's volatility is 172% a year, and its Sharpe ratio, 0.25, is well below the market's 0.48. In its worst month the portfolio lost 385% of its value, so an investor who followed it would have been wiped out. The statistics in the table treat each month's return as if the investor could start over.

Adjusting the forecasts with (15) fixes the problem in the most direct way. Each January we estimate the slope of realized on forecast alphas from all earlier years. It is never positive, so the adjusted forecasts are zero, and the optimal portfolio is the market. That is the right answer. Past alphas carry no information about future alphas, and an investor who has no better forecasts should hold the market. The Treynor-Black model pays off only when the analysts' forecasts have some accuracy, and the adjustment is what ensures that the portfolio bets only as much as that accuracy justifies.

Practice Problems

These problems give you a chance to practice the concepts introduced in this chapter. Try to solve each one on your own before expanding the solution.

Solutions to all problems can be found at https://lorenzonaranjo.com/class-materials/investments/.

Unless stated otherwise, the market has an expected excess return of $\mathrm{E}(R_M) = 8\%$ and a volatility of $\sigma_M = 20\%$, so its Sharpe ratio is 0.40.

Problem 1 (Adding One Stock to the Market). Olivia Carter, a portfolio manager at Lakemont Investors, holds the market and believes Verdant Foods (VDF) has an alpha of 3%, a beta of 1.2, and a residual volatility of 30%.

a. Use (9) to compute λ, and find the weights in VDF and in the market.
b. Compute the beta, expected excess return, and volatility of the optimal portfolio, and its Sharpe ratio.
c. Check the Sharpe ratio with (14).

Problem 2 (Profiting from a Negative Alpha). Olivia's analysts now conclude that Brackley Retail (BRK) is overpriced: it has an alpha of -2%, a beta of 0.9, and a residual volatility of 25%. Consider a portfolio of BRK and the market only.

a. Find λ, the weight in BRK, and the weight in the market.
b. Does BRK raise or lower the Sharpe ratio of the portfolio? By how much?

Problem 3 (A Two-Stock Active Portfolio). Nathan Price at Westbury Capital combines the market with two stocks: Arlo Systems (ARS), with an alpha of 2%, a beta of 1.2, and a residual volatility of 25%; and Belmont Paper (BMP), with an alpha of -1%, a beta of 0.8, and a residual volatility of 20%.

a. Compute λ and the weights in ARS, BMP, and the market.
b. Compute the beta of the optimal portfolio in two ways, using (7) and using the weights.
c. Compute the Sharpe ratio of the optimal portfolio.

Problem 4 (Same Alpha, Different Risk). Two analysts at Westbury Capital each pitch a stock with an alpha of 2%. Cresswell Labs (CRL) has a residual volatility of 10%, and Duncan Exploration (DXP) has a residual volatility of 40%.

a. How do the optimal weights of the two stocks compare?
b. By how much does each stock raise the squared Sharpe ratio of the optimal portfolio?

Problem 5 (How Many Ideas Does a Fund Need?). Harriet Lowe is launching a fund that holds the market plus an active portfolio. Her research team can find stocks with

information ratios of 0.05 in absolute value, independent of each other. How many such stocks does she need for the fund's Sharpe ratio to reach 0.50?

Problem 6 (The Cost of False Alphas). Return to Nathan's portfolio in Problem 3. Suppose his analysts were wrong, and the true alphas of ARS and BMP are both zero.

a. Using the same weights, what are the portfolio's expected excess return and volatility?
b. Compare its Sharpe ratio with that of the market. What does the comparison say about acting on alpha forecasts?

Problem 7 (Monthly and Annual Information Ratios). A stock has an alpha of 0.25% per month and a residual volatility of 5% per month, with residuals independent from month to month. Compute its monthly and its annual information ratio.

Problem 8 (Choosing Between Two Managers). The Corliss Foundation holds the market and wants to add one market-neutral manager. Manager Xu promises an alpha of 4% with a residual volatility of 20%. Manager Yates promises an alpha of 2% with a residual volatility of 5%.

a. Compute each manager's information ratio, and the Sharpe ratio the foundation could reach by combining the market with each manager optimally.
b. A trustee prefers Xu because his alpha is twice as large. How would you respond?

Problem 9 (What Does a Sharpe Ratio Reveal?). A fund that holds the market plus an optimal active portfolio reports a Sharpe ratio of 0.55. What is the information ratio of its active portfolio, assuming the market's Sharpe ratio is 0.40?

Problem 10 (Two Steps for Two Stocks). Solve Problem 3 again in two steps.

a. Use (10) and (11) to find the weights of ARS and BMP in the active portfolio A, and its alpha, beta and residual variance.
b. Use (12) to find the weight of A in the optimal portfolio, and check the weights of ARS, BMP and the market against Problem 3.
c. Compute the information ratio of A and check it against the information ratios of the two stocks.

Problem 11 (Adjusting a Forecast). An analyst forecasts an alpha of 4% for a stock with a beta of 1 and a residual volatility of 20%. A regression of the alphas realized on the analyst's past forecasts has a slope of 0.25.

a. What is the adjusted alpha forecast?
b. Find the optimal weight in the stock with the raw and with the adjusted forecast.
c. Suppose the adjusted forecast is the true alpha. Compute the Sharpe ratio of each portfolio in (b) and compare them with the market's.

References

Treynor, Jack L., and Fischer Black. 1973. "How to Use Security Analysis to Improve Portfolio Selection." *Journal of Business* 46 (1): 66–86.