Introduction

One of the disadvantages of the Markowitz model is that it requires estimates of all covariances and expected returns. Estimating the covariance between two assets is not hard, and can be done with precision, but estimating the pairwise covariances of five hundred different assets is a large task. Indeed, for five hundred stocks we would have to estimate $500 \times 501/2 = 125{,}250$ different variances and covariances. One way to circumvent this problem is to realize that stock returns exhibit a factor structure driven by the market.

Indeed, for many stocks the market explains a large fraction of their variability. Whatever is not explained by the market is firm-specific risk. The index model formalizes this intuition by splitting the variance of each stock into systematic and idiosyncratic variance.

The Model

The single index model is a linear regression between the excess returns of a stock and the excess returns of the market portfolio. Let's denote by r_i the return of stock i over a given period, and define $R_i = r_i - r_f$ as the excess return over the risk-free asset. Similarly, denote by R_M the excess return of the market over the risk-free rate. The index model postulates

$$R_i = \alpha_i + \beta_i R_M + e_i, \tag{1}$$

where $\mathrm{E}(e_i) = 0$ and $\mathrm{Cov}(R_M, e_i) = 0$.

In the single index model, the beta with the market captures the exposure of all securities to systematic risk, which is the risk shared by all securities. Therefore, the error term is idiosyncratic. The single index model assumes that $\mathrm{Cov}(e_i, e_j) = 0$ for any two securities i and j whose returns are not perfectly correlated.

Variance Decomposition

The assumptions of the single index model imply that the variance of R_i can be split into two parts:

$$\begin{aligned}\sigma_i^2 &= \mathrm{V}(\alpha_i + \beta_i R_M + e_i) \\ &= \beta_i^2 \,\mathrm{V}(R_M) + \mathrm{V}(e_i) + 2\beta_i \,\mathrm{Cov}(R_M, e_i) \\ &= \beta_i^2 \sigma_M^2 + \sigma^2(e_i).\end{aligned} \tag{2}$$

The first component of σ_i^2 is the systematic variance, which depends on the beta of the security but also the variance of the market. The second term of the variance of i is typically computed as the difference between the variance of i and the systematic variance.

Example 1 (Decomposing the Variance)**.** Suppose that you have the following regression for stock A:

$$R_A = \alpha_A + \beta_A R_M + e_A,$$

where $\alpha_A = 0.02$, $\beta_A = 1.2$, $\sigma(e_A) = 30\%$ and $\sigma_M = 25\%$. The variance of A can be computed as follows:

$$\sigma_A^2 = 1.2^2 \times 0.25^2 + 0.30^2 = 0.09 + 0.09 = 0.18.$$

In the previous expression, the systematic and idiosyncratic variances are the same. The standard deviation of A is then $\sigma_A = \sqrt{0.18} = 42.43\%$. □

The Security Characteristic Line (SCL)

The beta of the model can be estimated from the covariance of R_i and R_M, and the variance of R_M. Indeed, we have that

$$\begin{aligned}\mathrm{Cov}(R_i, R_M) &= \mathrm{Cov}(\beta_i R_M, R_M) + \mathrm{Cov}(e_i, R_M) \\ &= \beta_i \,\mathrm{V}(R_M),\end{aligned}$$

where we use the fact that $\text{Cov}(R_M, R_M) = \text{V}(R_M)$. Thus, in the index model we must have that

$$\beta_i = \frac{\text{Cov}(R_i, R_M)}{\text{V}(R_M)}. \tag{3}$$

The alpha of the security can then be computed as

$$\alpha_i = \text{E}(R_i) - \beta_i \, \text{E}(R_M). \tag{4}$$

The line

$$y = \alpha_i + \beta_i x$$

is called the *security characteristic line* (SCL) of security i. If we plot this line using a line chart, α_i is the intercept and β_i is the slope coefficient of the line.

R-Squared

Following the statistics literature, the proportion of systematic variance to total variance is called the R-squared of security i and can be expressed as

$$\text{R-squared} = \frac{\beta_i^2 \sigma_M^2}{\sigma_i^2} = 1 - \frac{\sigma^2(e_i)}{\sigma_i^2}. \tag{5}$$

Therefore, the R-squared can also be expressed as one minus the proportion of idiosyncratic variance to total variance. Since the single index model aims to decompose the total variance of a security into two orthogonal components, the R-squared gives us the proportions of this decomposition.

Example 2 (From R-Squared to Volatility)**.** You regress the excess returns of stock B on the excess returns of the market:

$$R_B = \alpha_B + \beta_B R_M + e_B.$$

Your regression package reports that $\alpha_B = -0.01, \beta_B = 0.8$ and the R-squared is 0.4. If the volatility of the market is 25% per year, the systematic variance is $0.8^2 \times 0.25^2 = 0.04$. Since

40% of the variance is systematic, we have that

$$\sigma_B^2 = \frac{0.04}{0.4} = 0.10,$$

which implies that $\sigma_B = \sqrt{0.10} = 31.62\%$ per year. We also know that 60% of the variance is firm-specific, which means that

$$\sigma^2(e_B) = 0.6 \times 0.10 = 0.06,$$

so that $\sigma(e_B) = \sqrt{0.06} = 24.49\%$ per year. □

Equation (5) can also be expressed in terms of the correlation between R_i and R_M. Indeed, since

$$\beta_i = \frac{\text{Cov}(R_i, R_M)}{\text{V}(R_M)} = \frac{\sigma_i \sigma_M \rho_{i,M}}{\sigma_M^2} = \frac{\sigma_i \rho_{i,M}}{\sigma_M}, \tag{6}$$

we have that

$$\text{R-squared} = \frac{\beta_i^2 \sigma_M^2}{\sigma_i^2} = \frac{\frac{\sigma_i^2 \rho_{i,M}^2}{\sigma_M^2} \sigma_M^2}{\sigma_i^2} = \rho_{i,M}^2. \tag{7}$$

Thus, the R-squared of a regression of R_i on R_M is just the square of the correlation between R_i and R_M. The name R-squared comes from the fact that we typically use the Greek letter *rho* (ρ), which corresponds to the Latin letter *r*, for correlation.

An Example with Real Data

As an example, let's analyze the monthly returns of Microsoft (Ticker: MSFT) from June 2015 until June 2025. Stock prices come from Yahoo Finance. As a proxy for the risk-free rate, I use the 3-month Treasury bill rate from FRED (Series: DTB3). The rate is expressed per year, so I convert it to a monthly rate by:

$$r_{\text{monthly}} = (1 + r_{\text{annual}})^{1/12} - 1.$$

To proxy for the market, I use the SPDR S&P 500 ETF Trust (Ticker: SPY), which allows me to include the dividend distribution of the stocks forming the S&P 500. The monthly returns

are computed using the *adjusted price* series to obtain a holding period return that includes dividends.

The table below presents some descriptive statistics of both series of excess monthly returns.

Table 1: The table presents descriptive statistics of Microsoft and S&P 500 monthly returns using data from June 2015 until June 2025.

Ticker	MSFT	SPY
Mean (%)	2.173	1.000
St. Dev. (%)	6.049	4.465

The table shows that during the period, the monthly returns of Microsoft are more volatile than those of the S&P 500, and so is its average excess return.

The figure below presents a scatter plot of the data. Clearly, the points cluster around the SCL, and we can see that the range of returns of Microsoft is significantly wider than the S&P 500.

We can use *ordinary least squares* (OLS) to estimate the slope coefficient and the intercept of the SCL. Many statistical packages can do this. The results below are computed using the Python library *statsmodels*.

```
                            OLS Regression Results
==============================================================================
Dep. Variable:                   MSFT   R-squared:                       0.522
Model:                            OLS   Adj. R-squared:                  0.518
No. Observations:                 120   F-statistic:                     114.5
Covariance Type:                  HC1   Prob (F-statistic):           4.29e-
19
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
Intercept      0.0120      0.004      3.194      0.001       0.005       0.019
SPY            0.9784      0.091     10.703      0.000       0.799       1.158
==============================================================================

Notes:
[1] Standard Errors are heteroscedasticity robust (HC1)
```

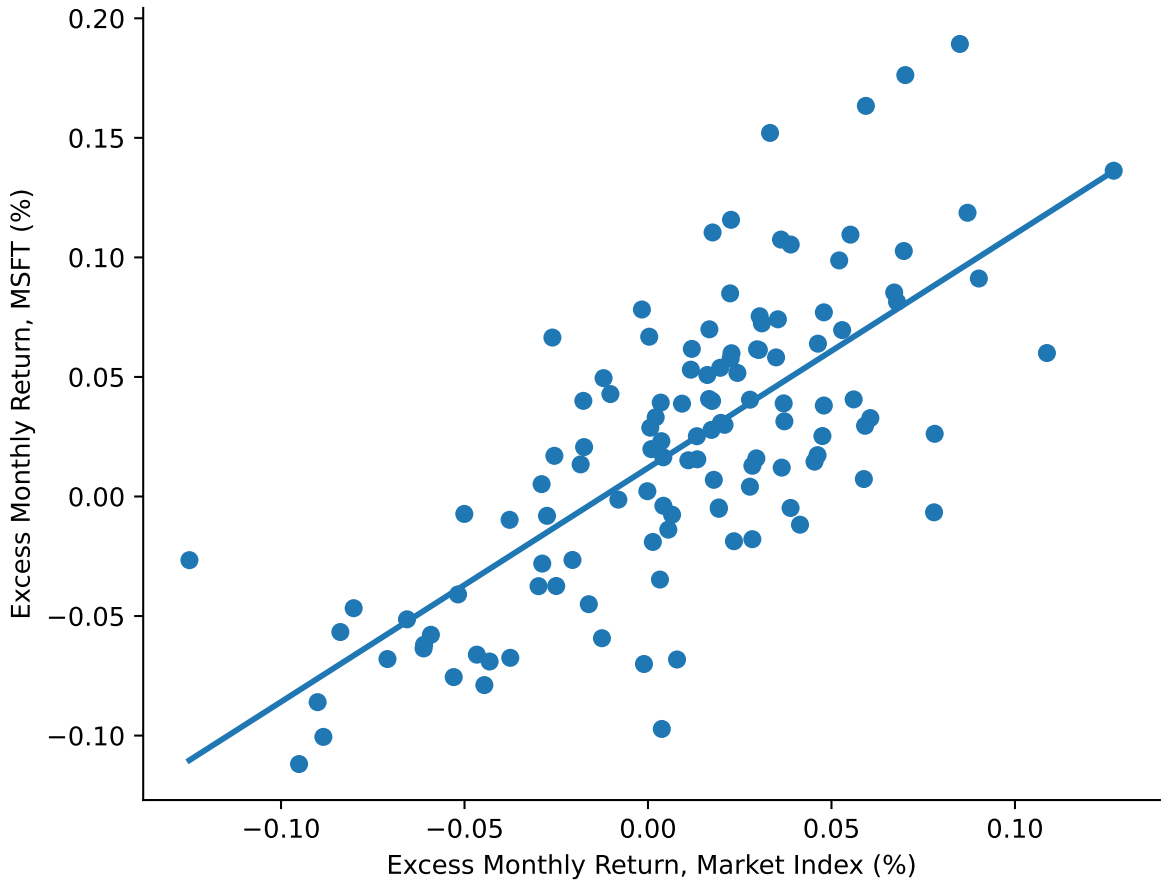


Figure 1: The figure plots the excess monthly returns of the market portfolio, proxied by SPY vs. the excess monthly returns of MSFT from June-2015 until June-2025.

The table presents many numbers but for our purposes we can focus on just a few of them. First, under the column `coef` we can see that the estimate for the intercept is 0.0120 whereas the beta estimate for MSFT is 0.9784. Since the beta of Microsoft is close to one, we can conclude that Microsoft carries almost the same systematic risk as the market.

We can compute the R-squared of the regression using the beta and variance of MSFT, and the variance of the market:

$$\text{R-squared} = \frac{0.9784^2 \times 4.465^2}{6.049^2} = 0.522.$$

The variance of MSFT and the S&P 500 are computed by squaring the standard deviations reported in Table 1. We can see that the computed R-squared corresponds to the R-squared reported by the regression package, implying that about 52% of the variance is explained by the exposure of MSFT to the market. The remaining variance is firm-specific risk.

Differences in Beta and R-Squared

The figure below shows a scatter plot of market excess monthly returns vs. the monthly excess returns of two financial and two technology stocks:

- Citigroup (Ticker: C)
- BlackRock (Ticker: BLK)
- Nvidia (Ticker: NVDA)
- Tesla (Ticker: TSLA)

As before, we proxy the market portfolio using the SPDR S&P 500 ETF Trust (Ticker: SPY), which allows us to include the dividend distribution of the stocks forming the S&P 500. The risk-free rate is obtained from the 3-month Treasury bill rate from FRED.

The picture shows that different stocks have different degrees of firm-specific risk. BlackRock excess returns align quite well with market returns, whereas Tesla excess returns are the most

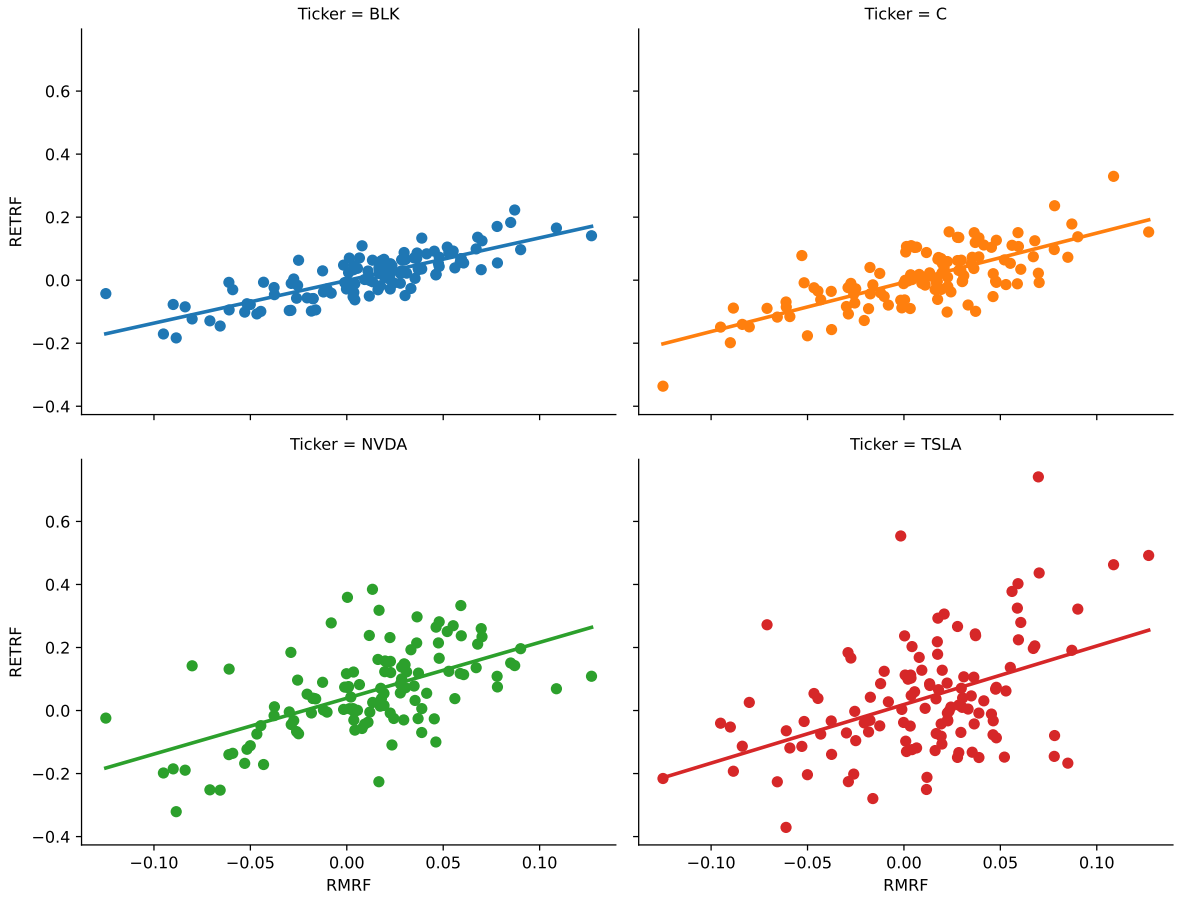


Figure 2: The figure plots the excess monthly returns of the market portfolio labeled as RMRF and proxied by SPY, vs. the excess monthly returns of Citigroup (C), BlackRock (BLK), Nvidia (NVDA), and Tesla (TSLA) labeled RETRF, from June-2015 until June-2025.

dispersed. In general, you would expect more dispersion in technology stocks since by definition new technologies might or might not work. The fact that you get a technology to work is independent of what the market does.

The table below presents the relevant values of the regression.

Table 2: The table presents the Alpha, Beta and R-Squared estimates obtained by running a linear regression of excess monthly returns for Citigroup (C), BlackRock (BLK), Nvidia (NVDA), and Tesla (TSLA) on the market portfolio proxied by SPY using data from June-2015 until June-2025.

	Alpha		Beta	R-Squared
	Estimate	P-value		
BLK	-0.0010	0.805	1.353	0.668
C	-0.0068	0.240	1.562	0.550
NVDA	0.0390	0.000	1.772	0.337
TSLA	0.0190	0.182	1.859	0.206

First, we can see that the R-squared is the lowest for Tesla, which is apparent from the pictures. The market explains a small fraction of the variance for the stock. On the other hand, about 67% of BlackRock's variance is explained by the market.

Also, we see that both Nvidia and Tesla load on significant systematic risk. Certainly, their cash flows are exposed to how the economy does and this is reflected on their high betas. Citigroup also has a high beta, which is typical of financial firms that are also exposed to how the market performs. BlackRock has the lowest beta of the four stocks, although it is still higher than the beta we estimated for Microsoft.

During the period the only stock that has a positive alpha statistically different from zero at the 1% significance level is Nvidia. For the other stocks, we cannot reject the null hypothesis that the alpha is zero. The alpha of Nvidia is indeed impressive. It has out-performed the market by 3.90% per month during the last 10 years.

Finally, the table below shows the pairwise correlations between each stock and the market.

Table 3: The table shows the pairwise correlations and the square of their values between excess monthly returns for Citigroup (C), BlackRock (BLK), Nvidia (NVDA), and Tesla (TSLA) and the market portfolio proxied by SPY using data from June-2015 until June-2025.

	Correlation	Correlation Squared
BLK	0.818	0.668
C	0.741	0.550
NVDA	0.580	0.337
TSLA	0.454	0.206

As expected, the square of the correlation corresponds to the R-squared reported in Table 2.

Not All Alphas Are Equal

Nvidia's alpha is large, but an alpha is only worth as much as the risk you have to bear to collect it. Buying a stock to earn its alpha also means holding its firm-specific risk e_i, and the market pays nothing for bearing that risk. Two stocks with the same alpha can therefore be very different opportunities.

Example 3 (A Lottery Ticket at Half Price)**.** Suppose a promoter sells lottery tickets at half their expected payoff: each ticket costs \$1 and pays \$1,000,000 with probability two in a million. The expected payoff is \$2, so the ticket's expected return is 100%, and since the draw has nothing to do with the market, its beta is zero and its alpha is essentially 100%. The payoff has a standard deviation of

$$10^6 \times \sqrt{2 \times 10^{-6} \times (1 - 2 \times 10^{-6})} \approx \$1{,}414,$$

so on a \$1 price the volatility of the ticket's return is about $141{,}400\%$, all of it firm-specific. Compare a stock with an alpha of 2% and a residual volatility of 10%. Its alpha is fifty times smaller, but it comes with far less risk. The discount does not justify a large position in tickets: a bigger bet on a single draw only scales up a risk that nobody compensates. □

The right comparison scales each alpha by the firm-specific risk that comes with it. The **information ratio** of security i is

$$IR_i = \frac{\alpha_i}{\sigma(e_i)}. \tag{8}$$

The lottery ticket has an information ratio of $1/1{,}414 \approx 0.0007$, while the stock's is $0.02/0.10 = 0.2$. The Treynor-Black model shows that the optimal position in each security is proportional to $\alpha_i/\sigma^2(e_i)$, which is 2 for the stock and $1/1{,}414^2 = 0.0000005$ for the ticket. The stock deserves about four million times the weight.

What changes the picture is diversification. A million tickets on independent draws keep the alpha of 100%, but the volatility of the average payoff falls by a factor of $\sqrt{10^6} = 1{,}000$, to about 141%. The information ratio rises to about 0.7, better than the stock's. Only when the firm-specific risk can be diversified away does a discount become a reliable profit. Real lotteries work the other way around: tickets sell above their expected payoff, and the operator, who sells millions of them, collects that margin almost without risk.

Implications of the Model

When you run a regression of the excess returns of a security on the excess returns of the market, the residuals are automatically orthogonal to the regressor. Therefore, in equation (1) we must have that $\text{Cov}(R_M, e) = 0$ for all securities.

The crucial assumption of the single index model is that the only systematic source of risk is the exposure of each security to the market. The implication of this assumption is that the covariance of the residuals between two securities is zero as long as their excess returns are not perfectly correlated.

We saw before that if A and B are perfectly correlated, we must have that $R_B = wR_A$ for some $w \neq 0$. If this were the case, the covariance between e_A and e_B is not zero even though they are different assets. More precisely, we have that

$$\text{Cov}(e_A, e_B) = \text{Cov}(e_A, we_A) = w\sigma^2(e_A) \neq 0,$$

provided that $w \neq 0$, i.e., you do not invest everything in the risk-free asset.

In the following, when we talk about two different assets it is implicitly assumed that their excess returns are not perfectly correlated, unless stated otherwise.

Covariance Structure

The covariance of the excess returns between securities i and j is

$$\begin{aligned}\text{Cov}(R_i, R_j) &= \text{Cov}(\alpha_i + \beta_i R_M + e_i, \alpha_j + \beta_j R_M + e_j) \\ &= \beta_i \beta_j \,\text{Cov}(R_M, R_M) \\ &= \beta_i \beta_j \sigma_M^2.\end{aligned} \tag{9}$$

The previous expression says that in the single index model, the covariance between any two different securities is given by their exposures to the market and the variance of the market. We can express equation (9) in terms of correlations as

$$\begin{aligned}\rho_{i,j} &= \frac{\text{Cov}(R_i, R_j)}{\sigma_i \sigma_j} = \frac{\beta_i \beta_j \sigma_M^2}{\sigma_i \sigma_j} \\ &= \frac{\frac{\sigma_i \rho_{i,M}}{\sigma_M} \frac{\sigma_j \rho_{j,M}}{\sigma_M} \sigma_M^2}{\sigma_i \sigma_j} \\ &= \rho_{i,M} \rho_{j,M}.\end{aligned} \tag{10}$$

Therefore, in the single-index model, all pairwise correlations between two assets can be computed as the product of their correlations with the market.

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/.

Problem 1 (Counting the Inputs)**.** Ingrid Halvorsen, a quantitative analyst at Halcyon Asset Management, covers a universe of 500 stocks and wants to build mean-variance efficient portfolios from it.

a. How many estimates does she need if she works directly with expected returns, variances, and covariances?
b. How many does she need under the single index model?

Problem 2 (Reading a Regression)**.** A regression of Quillon Robotics' (QLR) excess returns on the market's excess returns gives a beta of 1.5 and a residual volatility of $\sigma(e_{QLR}) = 35\%$ a year. The market's volatility is 18%.

a. Compute QLR's systematic variance, total variance, and volatility.
b. What R-squared should the regression report?

Problem 3 (Covariances and Portfolios Through Betas)**.** Two stocks in Lena Fischer's portfolio have the following index model estimates. The market's volatility is 20%.

Stock	Beta	Volatility
Corbel Aerospace (CBA)	1.2	40%
Dunmore Utilities (DMU)	0.7	25%

a. Compute the covariance and the correlation between the two stocks implied by the index model. Check the correlation using each stock's correlation with the market.
b. Compute each stock's residual variance.
c. Lena holds 60% in CBA and 40% in DMU. Compute the portfolio's beta and volatility using the index model.

Problem 4 (Building a Covariance Matrix)**.** Aisha Rahman needs a covariance matrix for three stocks to feed into a portfolio optimizer. She uses the index model with a market volatility of 20% and the following estimates:

Stock	Beta	Residual Volatility
Fernhill Grocers (FHG)	0.8	20%
Gallant Media (GLM)	1.0	25%
Hexley Mining (HXM)	1.5	30%

Compute the full covariance matrix of the three stocks' excess returns.

Problem 5 (How Many Stocks Are Enough?)**.** Marcus Bell builds equally weighted portfolios from a universe of stocks with an average beta of 1.1 and a residual volatility of 40% each. The market's volatility is 20%, and the index model holds.

a. Compute the volatility of an equally weighted portfolio of 10 stocks and of 50 stocks.
b. What volatility does the portfolio approach as the number of stocks grows without bound?
c. How many stocks does Marcus need for the portfolio's volatility to be within one percentage point of that limit?

Problem 6 (Estimating a Beta by Hand)**.** Sam Okoro, an intern at Brightline Advisors, has five months of excess returns, in percent, for Tamarind Foods (TMF) and the market:

Month	Market	TMF
1	3	3
2	-1	-3
3	5	9
4	-3	-3
5	1	4

Use (3) and (4) to estimate TMF's beta and alpha, and compute the R-squared of the regression.

Problem 7 (Is the Alpha Real?)**.** A regression of Elmstead Pharmaceuticals' (EPH) monthly excess returns on the market's gives an alpha of 0.5% per month with a p-value of 0.30, and a beta of 1.3. The market's expected excess return is 0.6% per month.

a. Using the point estimates, what is EPH's expected excess return per month?
b. A colleague, Ruth Abernathy, says EPH beats the market by half a percent a month. How would you respond?
c. Suppose the alpha is real. Fairmont Devices (FMD) has the same alpha of 0.5% per month, but a residual volatility of 4% per month against EPH's 12%. Which alpha is more valuable? How should the positions in the two stocks compare?

Problem 8 (When the Residuals Are Correlated)**.** Two regional banks, Harrow Bank (HRW) and Penfold Bank (PNF), have correlations with the market of 0.7 and 0.5. Over the same period, the sample correlation between their excess returns is 0.6.

a. What correlation does the single index model predict?
b. What does the gap say about the residuals of the two banks, and why might it arise?
c. Would the index model overstate or understate the risk of a portfolio made only of bank stocks?

Problem 9 (Scratch Cards at Half Price)**.** Rowan Gaming runs a promotion that sells scratch cards for $2. Each card pays $500 with probability $1/125$ and nothing otherwise, independently of other cards and of the market. Carla Mendoza is deciding whether the promotion is a better opportunity than Tidewell Shipping (TWS), a stock with an alpha of 3% and a residual volatility of 15%. Ignore the risk-free rate over the life of the promotion.

a. Compute the expected return, the volatility, and the information ratio of a single card.
b. Compare the card's information ratio with that of TWS. If positions are proportional to $\alpha_i/\sigma^2(e_i)$, how does the position in one card compare with the position in TWS?
c. Carla can instead spread the same dollars over N cards with independent outcomes. How many cards does she need for the cards to match the information ratio of TWS?

Problem 10 (Hedging Out Market Risk)**.** Garnet Capital, a hedge fund, buys $10 million of Lumen Optics (LMO). Its analysts estimate an alpha of 2% a year, a beta of 1.4, and a residual volatility of 25% a year. The fund's manager, Tomas Varga, wants a bet on LMO's alpha alone, without exposure to the market.

a. How large a short position in an S&P 500 index fund should Tomas take?

b. What are the expected excess dollar return and the dollar volatility of the hedged position over a year?