Pearl Fund
Portfolio Optimization & Statistical Analysis

Navein Suresh — Suryaveer Singh

March 20, 2024

Motivation

The primary objective of this project is to enhance our proficiency in the R programming language while applying key statistical and financial concepts acquired in previous coursework, including Asset Pricing and Portfolio Theory, Econometrics, Empirical Asset Pricing, and Advanced Programming in R. Through the development of the Pearl Fund, a diversified portfolio of selected stocks and ETFs, we aim to demonstrate our ability to conduct portfolio optimization and analysis using R.

If you would like to test your own curated portfolio of assets, feel free to download our code and enter the exchange-traded tickers of your assets into the "stock_ticker" variable at the very beginning. The only caveat is that any company you enter must have gone public (exist on a public exchange) since at least 2004.

Fund Details

The Pearl Fund is a portfolio designed to achieve a balance between diversification and stability in returns. This selection includes a mix of sectors such as technology, consumer goods, healthcare, financial services, energy, and defense, ensuring that the portfolio is well-diversified and less susceptible to sector-specific risks.

Companies like Microsoft and HP Inc. represent the technology sector, known for its growth potential, while consumer giants like Costco, PepsiCo, and Procter & Gamble offer stable returns and resilience in various economic conditions. The inclusion of healthcare stalwart Johnson & Johnson adds a layer of defensive positioning, given the inelastic demand for healthcare products and services.

The financial sector is represented by JPMorgan Chase, providing exposure to the banking industry and its cyclical nature. Energy sector exposure is achieved through Chevron Corporation, which can offer hedge against inflation and benefit from rising energy prices. Raytheon Technologies brings in the defense angle, a sector known for its steady demand irrespective of economic cycles. The addition of Ford Motor Company gives the portfolio a touch of the cyclical automotive industry, which can provide growth in economic upswings.

Lastly, the inclusion of the iShares 20+ Year Treasury Bond ETF (TLT) adds a fixed income component to the portfolio, offering stability and income, which is crucial for reducing overall portfolio volatility.

Pearl Fund's diversity allows for the application of various optimization techniques to achieve an efficient risk-return trade-off.

Benchmarks

The S&P 500 and the NASDAQ serve as effective benchmarks for the Pearl Fund due to their broad representation of the market and relevance to the fund's holdings. The S&P 500, comprising 500 of the largest U.S. companies across various sectors, provides a benchmark for the overall market performance and is particularly relevant given the Pearl Fund's inclusion of large-cap stocks from diverse industries.

The NASDAQ, known for its heavy technology sector weighting, is a suitable benchmark for assessing the performance of the technology components within the Pearl Fund, such as Microsoft and HP Inc. By comparing the Pearl Fund's performance against these benchmarks, we can evaluate its relative success in capturing market trends and generating returns above these well-established indices.

Financial Theory

When looking through our analysis, its crucial to have an understanding of the financial and statistical underpinnings. We've included a summary of the one's you should familiarize yourself with below.

Variance-Covariance Matrix: The variance-covariance matrix is a tool used in portfolio analysis to measure the volatility and correlation between assets. The diagonal elements represent the variance (risk) of each asset, while the off-diagonal elements show the covariance, indicating how the returns of different assets move together. This matrix is crucial for optimizing a portfolio's risk-return profile, as it helps in determining the combination of assets that minimizes risk for a given level of expected return.

Correlation Matrix: The correlation matrix is used in portfolio analysis to measure the degree to which the returns of different assets move together. Each element ranges from -1 to 1, with 1 indicating perfect positive correlation, -1 indicating perfect negative correlation, and 0 indicating no correlation. This matrix helps in understanding the diversification benefits within a portfolio and is crucial for asset allocation and risk management. For viewing purposes, we represent this as percentages.

Principal Component Analysis: Principal Component Analysis (PCA) is a statistical technique used to reduce the dimensionality of a dataset while retaining most of the variation in the data. It transforms the original variables into a new set of uncorrelated variables called principal components, which are ordered by the amount of variance they capture. The graph can help visualize how different assets in a portfolio group together or differ based on the underlying factors represented by the principal components. It can also show how much of the portfolio's variance is explained by each principal component, aiding in the interpretation of the data and in making informed investment decisions.

Correlogram: A correlogram is a graphical tool used to visualize the correlation coefficients between pairs of variables in a dataset. It typically displays a matrix with colors or shading to represent the strength and direction of correlations, with values ranging from -1 (perfect negative correlation) to +1 (perfect positive correlation).

Active Premium: The additional return that an active investment strategy generates compared to a benchmark index.

Alpha: A measure of the excess return of an investment relative to the return of the benchmark index.

Annualized Alpha: The average annual excess return of an investment compared to the benchmark.

Beta: A measure of the volatility, or systematic risk, of a security or portfolio in comparison to the market as a whole.

Beta- : A measure of downside beta, reflecting the sensitivity of an asset's returns to negative market movements.

Beta+ : A measure of upside beta, reflecting the sensitivity of an asset's returns to positive market movements.

Correlation p-value: The p-value associated with the statistical test of the correlation coefficient, indicating the probability that the observed correlation occurred by chance.

Information Ratio: A measure of the risk-adjusted return of an investment, calculated as the excess return of the investment relative to a benchmark, divided by the standard deviation of the excess returns. R-squared: A statistical measure that represents the proportion of the variance for a dependent variable that's explained by an independent variable or variables in a regression model.

Tracking Error: The standard deviation of the difference between the returns of an investment and its benchmark, used to measure the consistency of excess returns.

Treynor Ratio: A performance metric for determining how much excess return was generated for each unit of risk taken on by a portfolio, using beta as the risk measure.

Efficient Frontier: The efficient frontier is a curve representing a set of optimal portfolios that

offer the highest expected return for a given level of risk or the lowest risk for a given level of expected return. It is a key concept in modern portfolio theory, used to identify the best possible investment portfolios based on their risk-return profiles.

Expected Tail Loss: Also known as Conditional Value at Risk (CVaR) or Expected Shortfall, ETL is a risk measure that estimates the expected loss in the worst-case scenario of a given confidence level. It focuses on the tail end of the loss distribution, providing a more conservative estimate of risk than Value at Risk (VaR), which only indicates the maximum potential loss at a specific confidence level without considering the magnitude of losses beyond that point.

Optimization

For the process of optimization, we use the ROI package. The ROI package serves as an interface to the Rglpk package and the quadprog package to solve linear and quadratic programming problems. Our optimization process consists of defining our constraints & objectives. We then optimize our asset allocation weights using historical returns from 2004 till today. Other optimization solvers are available (check references) which may serve better depending on certain situations. Our main constraints include ensuring no leverage (borrowing), weight sum must add up to 1. Our core objective is to maximize returns while minimizing our variance under these constraints.

References

We would like to thank Dr. Stephen Bianchi (Professor of Economics @ U.C. Berkeley) & Dr. Gaston Sanchez (Professor of Statistics @ U.C. Berkeley) for their advice, guidance, and mentorship throughout this project. Some additional resources we consulted are listed below.

1. Basic Functions in Tidy Quant (https://cran.r-project.org/web/packages/tidyquant/vignettes/TQ01-core-functions-in-tidyquant.html)
2. Guide to Portfolio Optimization (https://miltonfmr.com/the-complete-guide-to-portfolio-optimization-in-r-part1/)