Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
May 2, 2024
A Regional Conference of the Society for Industrial and Systems Engineering

Potential Value Investing Strategy within the TSP

Caleb Kish and Nathan Hedgecock

Department of Systems Engineering,
United States Military Academy,
West Point, NY 10996

Corresponding author's Email: calebwork85@gmail.com

Author Note: Caleb Kish is a current student in the Department of Systems Engineering at West Point, particularly interested in investment and personal finance. A special thanks to MAJ Hedgecock for advising this work, LTC Pedersen for assisting with machine learning concepts and the model setup, and CDT Tobias Hild for being invaluable in expediting the data collection process. The views expressed herein are those of the authors and do not reflect the position of the United States Military Academy, the Department of the Army, or the Department of Defense.

Abstract: The Thrift Savings Plan (TSP) is the retirement investment account for government employees and military. The TSP functions similarly to a 401k or Roth IRA and provides a large portion of retirement income for investors. The TSP currently has five main funds that exhibit either low-risk, low-return (G, F) or high-risk, high-return (C, S, I). The primary objective of this research is to use value investing strategies to present a third viable option providing sufficient returns for moderate risk. Value investing strategies center around intrinsic company valuation, often based on income statements, balance sheets, cash flows, and other metrics. This requires significant amounts of data processing, which is meticulous and difficult without years of experience. This report explores the use of neural networks to process and predict data for investors. The model, while small, shows promise in predicting value investing variables through expansion and further data refinement.

Keywords: TSP, Value Investing, Machine Learning

1. Introduction

The Thrift Savings Plan (TSP) is the retirement investment account for all government employees and military personnel ("About the Thrift Savings Plan (TSP)"). Functioning similarly to a 401k or Roth IRA in the private sector, the TSP carries a significant portion of government employee's retirement portfolios. For this reason, it is important to give TSP investors as much control as possible over their accounts to best suit their preferences for their retirement portfolio. Based on high correlation between two main groups of TSP funds, value investing may present a feasible alternative to what is currently available. Value investing requires an estimation of the intrinsic value of a company based on metrics or company data. In order to process this vast amount of data, artificial intelligence in the form of a neural network was used. This model is presently aimed to predict the next quarter's change in stock price based on SEC report metrics shared between the companies in the dataset. This sets the stage for future improvements of the model to generate a sustainable TSP investment strategy based on value investing principles.

2. Literature Review

The TSP currently has five individual funds: G, C, F, S, and I. The G and F mirror government securities and bonds, while the C, S, and I mirror the S&P 500, DJIA, and MSCI EAFE (International), respectively ("Individual funds"). The TSP also offers a mutual fund window (MFW) that comes with a variety of restrictions and fees, making it difficult for most members to access (Carlos). Lastly, the TSP has a variety of Lifecycle (L) funds that combine individual funds based on investors' planned withdrawal date.

Initial analysis of historical TSP yields from the establishment of each fund, seen in Figures 1 and 2, shows that the G and F funds have historically exhibited more consistent lower yields while the C, S, and I have exhibited higher yields with significantly higher volatility ("Rates of return"). Figure 2 shows that investors seeking higher yields must expose themselves to the downside risk associated with the C, S, and I funds.

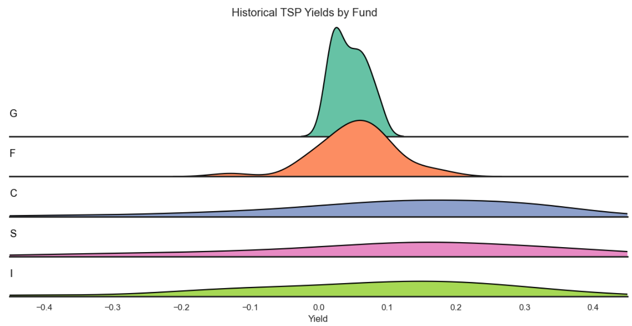

Figure 1. Historical TSP Yields by Fund (Histogram Curves)

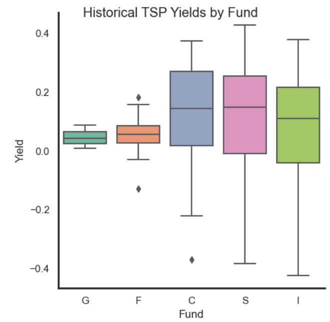

Figure 2. Historical TSP Yields by Fund (Boxplots)

While an L fund will provide some diversification and thus some protection against this risk, it should be noted that the base funds can exhibit high correlation, as seen in Figure 3, particularly during market downturns. Thus, it would be in the best interest of TSP investors to have access to a fund that would yield returns similar to the C, S, and I funds while exhibiting either less volatility or less correlation. This would allow L funds to mitigate risk better and give investors a wider variety of options to manage their retirement better, as seen in Figure 4.

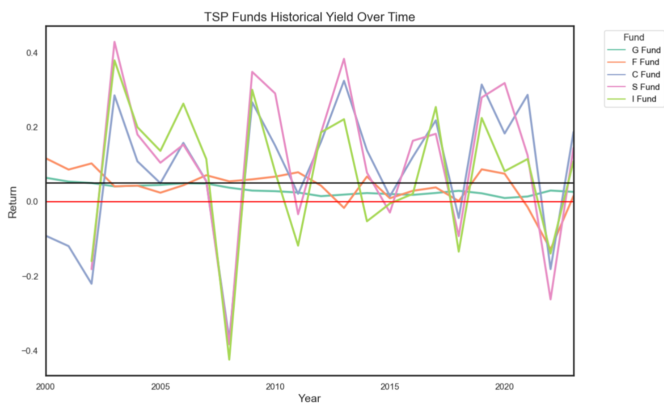

Figure 3. TSP Funds Historical Yields Over Time

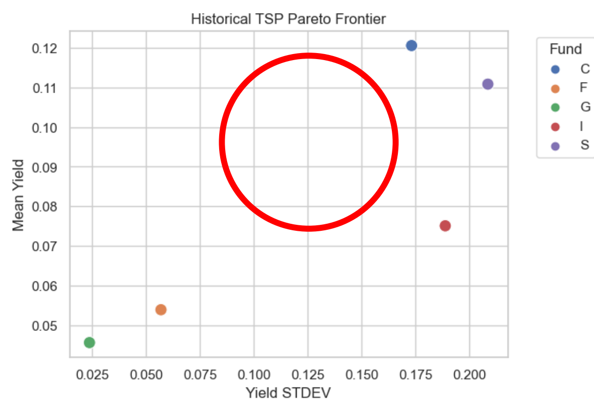

Figure 4. Historical TSP Yields by Fund (Pareto Frontier)

One factor that plays into the volatility of the C, S, and I funds is the types of securities present within them. Two common types of securities are growth and value. Growth stocks are often labeled and priced based on expectations of large future cash flows, while value stocks are generally seen as having more consistent cash flows (Gambarini). Since growth stocks involve a higher sense of speculation, they are often seen as more volatile. Value investing performs better during periods of high interest rates that inhibit companies' potential for aggressive growth, or during bearish markets when optimism is short (Gambarini, Cussen). This would provide valuable diversification for TSP investors who have historically been hit hard by major economic downturns. According to Cussen and Akilesh, the S&P contains about 29% value and 40% growth, while the Dow has about 50% value and 20% growth. This is likely a significant factor in TSP yields during major economic downturns, as seen in Figure 3. Thus, a value investing strategy would provide valuable diversification for TSP investors during market downturns or in the current economic climate when Federal Open Market Committee (FOMC) rates are high.

Value investing strategies generally revolve around buying companies at a price lower than their intrinsic value, based on the current valuation (rather than the future valuation, which would constitute a growth strategy). Investors often use metrics, exclusionary criteria, and published company data to value a company. Metrics are generally used to aid in company valuation or risk estimation, with some of the more common metrics in the Valuation and Risk Metrics section. Different investors combine different metrics and exclusion factors. These include Buffet's rules of at least 10 years in the market and avoiding technology sectors (The Investopedia Team), Greenblatt's Magic Formula with EBIT and Earnings Yield (Kenton, "Magic Formula Investing"), or Michael Burry's EV/EBITDA screen ("Michael Burry's Investing Formula", Jacobs). Regardless, all value investors delve into company data, particularly meeting minutes and reports publicly filed with the SEC.

SEC reports are long and monotonous; it often takes successful investors decades to develop an understanding of the data provided in the reports, hence the widespread use of summarizing metrics like P/E, PEG, etc. With the massive volume of data, it may be beneficial to leverage machine learning to parse patterns from the data rather than endlessly reading SEC reports. John Alberg and Lakshay Chauhan from Euclidean Technologies attempted to build a deep learning model to predict stock price, but given that price is subject to various factors, they were unable to build a successful model. However, they were able to adapt their model to predict future EPS. They did this to identify what they called "Value Traps" within a portfolio, defined as securities that ranked highly in EV/EBIT but settled in the bottom 10% of performers. This tool yielded them success in their portfolio simulations and backtests.

Another example is Abram Haich, who created two models to predict EPS and future cashflows. He used Compustat and Yahoo Finance data to pull together the fundamental data in income statements, balance sheets, and cashflow statements. It is important to note that a series of data manipulations, such as normalization and coefficient multiplication for stock splits, occurred prior to model training. His first model performed well predicting future EPS, specifically the spikes after the earnings statement was released. His discounted cash flow (DCF) model performed better than a linear model only in some instances. He believes a significant portion of his model's shortcomings came from major economic events and industry-specific factors not accounted for. According to backtests, the top five securities picked by his DCF model returned an average of 67% compared to the 33.38% returned by the S&P during the same periods (Haich 21-22). Haich notes his model could be improved by factoring in macroeconomic events, using more metrics than EPS and PE, and testing different techniques and architectures such as GANs or Transformers (Haich 28-29).

Overall, it seems that machine learning strategies have shown initial success but would need macroeconomic data to fully account for major market changes to optimize performance. This research will aim to build an initial machine learning model using SEC data combined with sector and price data from Yahoo Finance to build a baseline predictive model that will create future opportunities to mitigate these shortcomings.

3. Valuation and Risk Metrics

There are a variety of metrics used by value investors to aid in estimating the intrinsic value of a company. These metrics are most often based on accounting variables from quarterly and annual reports, as well as the current price of the security. For better understanding of the metrics listed above, a sample of important metrics are listed below.

P/E Ratio – Price to Earnings Ratio. Measures how much an investor pays for each dollar of company earnings. Calculated by dividing Price by Earnings Per Share (EPS).

Earnings Yield – Used as a gauge for optimal asset allocation (Mitchell). A higher Earnings Yield can signify an undervalued company. Calculated as the inverse of P/E ratio.

PEG – Price to Earnings Growth Ratio. Functions similar to P/E but considers the EPS growth rate during the last period (Kenton, "Price/Earnings-to-Growth"). A lower ratio is seen to signal better value. Calculated by dividing Price by EPS by EPS Growth Rate.

EV/EBITDA – Enterprise Multiple. A ratio between the value of a company's entire operations, calculated by adding Market Cap (Equity) and Debt then subtracting cash, and a company's Earnings Before Income Tax, Depreciation, and Amortization ("Michael Burry's Investing Formula"). This ratio can be used to compare companies in similar industries to see which are the most efficient at generating earnings based on the size of their operations.

4. Methodology

4.1 Data Acquisition

The aim of this project was to build a model to predict change in stock price based on company metrics. Data collection required pulling together common metrics shared by various companies' SEC filings to then merge with their respective price data. The company data was collected from the SEC EDGAR database containing all company facts sheets for historic and current companies. These files contained the information from every 10Q/10K filing for that company. Many companies add a variety of accounting components in addition to those required by the SEC, causing most files in the database to be formatted differently with a variety of unshared variables. Additionally, many common GAAP (General Accepted Accounting Principles) variables were not present wholistically, but rather broken up into less standardized subordinate variables. With over 5000 sheets available, the code used for JSON parsing could not accurately parse every fact sheet in the database, despite repeated revision. For this reason, 20 companies from the SP500 with well-parsed data were selected for the initial model. The data-cleaning algorithm eliminated another 8 fact sheets from this sample.

4.2 Stock Data Building

The data from each company was formatted with rows representing a quarterly filing and columns representing the accounting variables reported. Price data was pulled and merged from yfinance for each security at or around the filing date. This ultimately yielded a data frame for each company listing the reported accounting variables for each report as well as the corresponding price at the time of the report. Lastly, the variables in all data frames were filtered to contain only the variables shared amongst the 12 companies in their filings, approximately 15 variables in total not including price.

4.3 Model Data Building

For each company, the data frame was pivoted by combining four rows (quarters) into a single row to properly account for the time-series nature of the data (Pedersen). The output was then calculated as the next row's first adjusted close divided by the current row's last adjusted close. This ultimately created a data frame where each row contained the past four quarters of company data and the prices at each quarter, finally followed by a single output value representing the change in price from the most recent of the four quarters to the next quarter into the future. For example, one row may contain Q1-Q4 for FY2010 and the output would be the change in price from Q4 2010 to Q1 2011.

All dates were then converted to numbers using Excel, and columns containing more than 20 null values were dropped. Lastly, the output was run through a sigmoid function centered around 1 (x-1) and with an added constant C, to scale the change in price to between zero and one.

$$y = \frac{1}{1+e^{-C(x-1)}} \quad (1)$$

The center at one was chosen to center the function around no change in price. The constant was used to affect the functions sensitivity nearest to the center. Increasing the function shrinks the range of price changes that are not approximately 0 or 1. By using a constant of 10, the sigmoid function would classify any price changes more extreme than 14% loss or gain as 0.2 and 0.8 respectively. As well, it would classify price changes more extreme than 30% either way as approximately the same. Based on visual inspection, the function appears to fit well for the model, but alteration could have an affect on the accuracy.

4.4 Model Setup

The model chosen was a neural network due to the ease with which a neural network can account for interactions between variables. Based on the nature of the data and value investing principles, the aim is to generate a wholistic view of a company based on the metrics present in their quarterly and annual reports. A neural network is best suited for this task as it will account for the interactions between all variables and the interactions between interactions. These second and third order effects are incredibly important when gauging the health and potential future success of a company. For example, looking simply at the total value of assets is meaningless. However, accounting for depreciation provides a slightly better picture by providing context into the timescale and usability of said assets. Accounting for new assets added based on the previous quarter, particularly what amount of new assets are PPE, can be valuable for estimating current and future production capability. Future production capability would factor heavily into future cash flows, which are key in moving the price and dividend potential for a security. This is just a single example of how analyzing the intrinsic value of a company requires deep analysis of multiple variables and their interactions. Ultimately, a neural network is suitable as it can perform a similar analysis multiple times over, accounting for chains of interactions that may be missed even by human investors knowledgeable and seasoned in the analysis of business accounting reports.

The initial neural net was built using Adrian Tam's tutorial, starting with the variables common between the company fact sheets used, seen in red in Figure 5. Figure 5 was generated using Alexander Lenail's Neural Network SVG with each node representing 10 separate nodes for simplicity's sake, apart from the singular output. The initial model was incredibly basic, not splitting between a train and a test set. However, the model was able to get after the necessary interactions between variables by using three hidden layers with a node count of 150% of the variable count, followed by a final hidden layer with node count equal to the variable count as seen in blue. This model aimed to predict the change in price from the most recent of the last four quarters to the next quarter as seen in purple in Figure 5.

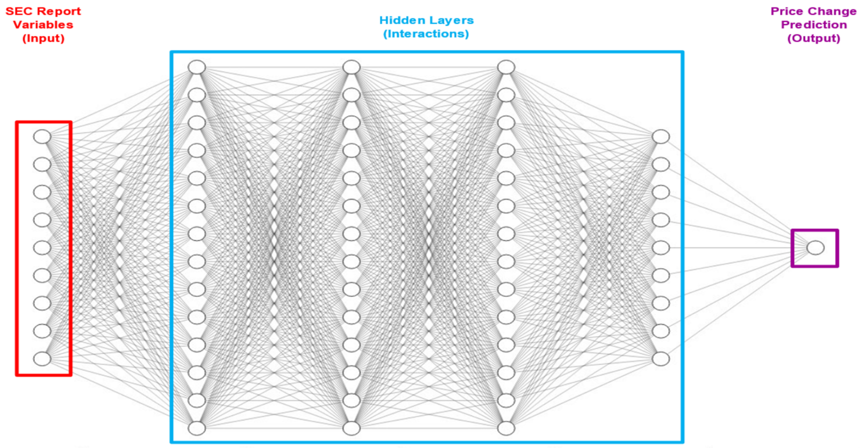

Figure 5. Neural Network Model (Visualized)

After following Tam's instructions for the base model, it was then iteratively improved using ChatGPT. First, the function within each node was changed from sigmoid to linear since change in price is a continuous variable rather than a classification. The loss function was then changed to MSE to properly suit a regression model. Additionally, the data was then split between a train (75%) and test (25%) set. After running the model, the output graphs seen in Figures 6 and 7 were generated.

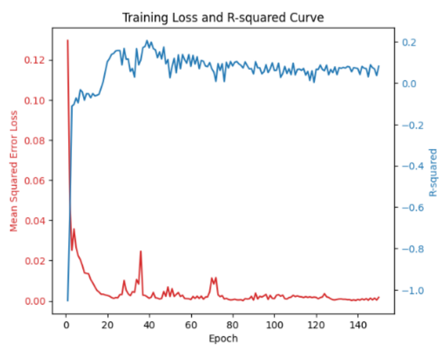

Figure 6. Training Loss and R^2 vs. Number of Epochs

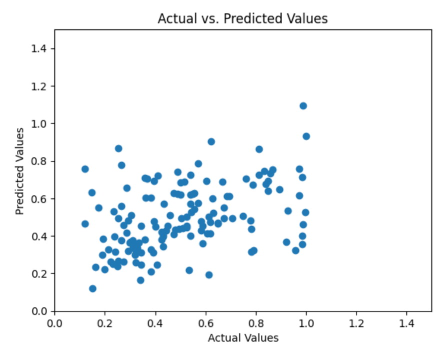

Figure 7. Actual vs. Predicted Values

All these changes led to a massive improvement from Tam's original model. The improved model was analyzed repeatedly using different amounts of hidden layers and different numbers of neurons within those layers. The random input for the test and train split was altered so the final result yielded a unique and unused split.

5. Results and Discussion

The MSE loss and R^2 by epoch count can be seen in Figure 5. The R^2 can be visualized in Actual vs. Predicted Values in Figure 6. The final MSE loss and R^2 values were 0.001207 and 0.0435, respectively. While these results were much less significant than desired, they yielded significant improvements compared to earlier trials containing almost 100 securities with nearly 500 variables. Also, there were significantly fewer outliers within the data compared to previous attempts. Overall, this points towards the possibility that using only variables shared amongst reports may significantly impact results. Despite shortcomings, the model shows promise for future improvements that could provide the first step towards a sustainable value investing strategy. At the minimum, tailoring the model towards predicting future metrics rather than price would decrease volatility and provide TSP investors with insights. This would be the start of a TSP strategy based on fundamentals. Alternatively, a better case would provide the TSP with a model that can either identify high yield investments, potential losses, or both. Nevertheless, the model is currently not accurate enough to be reliably integrated and needs significant additions before implementation.

ISBN: 97819384962-4-0

The largest limitation of this research was the ability to parse the data. Given the differences between quarterly reports, no singular method could be used for all available fact sheets. Future data parsing should focus on gathering specific variables or metrics such as EPS, EBITDA, P/B, etc. This would be significantly easier from sources outside the EDGAR database due to aforementioned parsing difficulties. Using widely accepted variables would improve data size and consistency and thus the model's overall usability. Lastly, the data was structured to consider the past four quarters. In terms of value investing, a one-year period is likely not enough time for a company's investments to gain fruition and affect share price.

The model could be improved by adding more securities for consideration within the data. The training process could be vastly improved by adding K-fold cross-validation, testing different constants for the sigmoid function, and testing different learning rates and optimization functions. Despite this model's shortcomings, it does provide a foundation for analyzing SEC report data using machine learning. Based on Figure 6, there is a significant possibility that with more data, the model could perform significantly with minimal tuning required. The output could then also be changed to predict cash flows, earnings, or other metrics deemed important to the investing decision-maker. Thus, future research will ultimately aim to provide tools for investors to more accurately value companies so they can be purchased at a discount. This is the first step necessary to create a value investing strategy within the TSP that will provide necessary diversification to investors.

6. References

"About the Thrift Savings Plan (TSP)." Thrift Savings Plan, https://www.tsp.gov/about-the-thrift-savings-plan-tsp/. Accessed 1 October 2023.

Alberg, John and Chauhan, Lakshay. "Identifying ValueTraps with Deep Learning." Euclidean Technologies, https://www.euclidean.com/value-traps-and-deep-learning. Accessed 5 October 2023.

Brock, Catherine. "What Is Value Investing And 4 Best Strategies." Forbes, 12 May 2023, https://www.forbes.com/sites/investor-hub/article/what-is-value-investing-and-4-best-strategies/?sh=55f44fc9732f.

Carlos, Alvin. "TSP Mutual Fund Window: What You Need To Know." District Capital Management, 3 June 2022, https://districtcapitalmanagement.com/tsp-mutual-fund-window/#:~:text=In%20June%202022%2C%20the%20Federal,to%20investors%20within%20the%20TSP.

Cussen, Mark. "Value or Growth Stocks: Which Is Better?" Investopedia, 27 July 2023, https://www.investopedia.com/articles/professionals/072415/value-or-growth-stocks-which-best.asp.

Gambarini, Simona. "GROWTH VS. VALUE: RETHINK YOUR INVESTMENT STYLE." Goldman Sachs, 20 April 2023, https://www.gsam.com/content/gsam/us/en/institutions/market-insights/gsam-insights/perspectives/2023/growth-vs-value-re-think-your-investment-style.html.

Haich, Abram. DEEP LEARNING APPLIED TO PUBLIC COMPANY VALUATION FOR VALUE INVESTING. 2021. North Dakota State University, master's Thesis. https://library.ndsu.edu/ir/bitstream/handle/10365/32668/Deep%20Learning%20Applied%20to%20Public%20Company%20Valuation%20for%20Value%20Investing.pdf?sequence=1&isAllowed=y.

"Individual funds." Thrift Savings Plan, https://www.tsp.gov/funds-individual/. Accessed 1 October 2023.

The Investopedia Team. "Warren Buffett's Investment Strategy." Investopedia, 18 September 2023, https://www.investopedia.com/articles/01/071801.asp.

Jacobs, Dillon. "A Fascinating Look at Dr. Michael Burry's Investment Strategy." Finmasters, 30 May 2022, https://finmasters.com/michael-burry-investment-strategy/#gref.

Kenton, Will. "Magic Formula Investing: Definition and What It Tells You." Investopedia, 8 July 2022, https://www.investopedia.com/terms/m/magic-formula-investing.asp.

Kenton, Will. "Price/Earnings-to-Growth (PEG) Ratio: What It Is and the Formula." Investopedia, 5 September 2022, https://www.investopedia.com/terms/p/pegratio.asp.

Lenail, Alexander. "NN-SVG." Accessed 17 April 2024. https://alexlenail.me/NN-SVG/index.html.

"Michael Burry's Investing Formula." YouTube, uploaded by Benjamin, 25 May 2021, https://www.youtube.com/watch?v=TrTeLB8F98k&ab_channel=Benjamin.

Mitchell, Cory. "Earnings Yield: Definition, Example, and How To Calculate It." Investopedia, 25 March 2022, https://www.investopedia.com/terms/e/earningsyield.asp.

Pedersen, LTC Joseph. Interview. Conducted by Caleb Kish. 06 November 2024.

"Rates of return." Thrift Savings Plan, https://www.tsp.gov/fund-performance/. Accessed 5 October 2023. This was used to collect historical data used in Figure 1 and 2.

Tam, Adrian. "Develop Your First Neural Network with PyTorch, Step by Step." Machine Learning Mastery, 8 April 2023, https://machinelearningmastery.com/develop-your-first-neural-network-with-pytorch-step-by-step/.