

S&P 500 Portfolio Optimization Using Macroeconomic Factor Models

David Newcomb

Mgmt. Science & Engineering
Stanford University

Zach Skokan

Mgmt. Science & Engineering
Stanford University

Thomas Stephens

Mgmt. Science & Engineering
Stanford University

Abstract: This paper seeks to examine the utility of macroeconomic factor models, which leverage observable economic data to measure and project stock returns. Our analysis focused on portfolio optimization within the S&P 500 universe, and accordingly focused on U.S. domestic factors. We used multiple approaches for factor selection: Akaike Information Criterion, Bayesian Information Criterion, and a hand-selected grouping of factors. These different selections—regressed from December 2003 to November 2008—suggest varying optimal portfolio investment strategies over the December 2008 to November 2013. Finally, we compared these results to a statistical factor model.

I. INTRODUCTION

When setting out on this study, we sought to implement a model, which used macroeconomic factors to predict out the S&P500 equities index to maximize portfolio return/minimize portfolio variance. But what is a factor model and why chose a macroeconomic one?

Factor models observe time series of ‘factors’ and then use that information to create predictive matrices on equity returns based off of the levels of said factors. These time series could be anything: from interest rates, to market cap, to the weather. Generally speaking, there are three main types of factor models with respect to security market returns: macroeconomic, fundamental, and statistical. According to Gregory Connor in his paper “The Three Types of Factor Models: A Comparison of Their Explanatory Power,” macroeconomic factor models have historically low predictive power as compared to their two contemporaries (*Financial Analysts Journal*, 42). A table specifying these differences, also courtesy of Connor, can be found in Appendix Figure A.

Despite their inferiority, macroeconomic factor models still serve multiple important purposes. First, they provide a learning opportunity for creating factor models in general, as the formulation for all of the models is for the large-part equivalent. Secondly, they are innately easier to understand. Many have

heard statements such as “Oil is going [here] so the market is going [there]” or similar statements. Macroeconomic factor models test the legitimacy of aforementioned statements. Lastly, macroeconomic factors provide a unique opportunity for factor selection, as a plethora of information regarding the macro economy exists. As one will see in part III of this study, this served as a sub-puzzle to the study in full. First, we will walk through our formulation of a factor model. Then we will move to our regression techniques, then to the optimization itself. Finally, we will finish with results and conclusions.

II. FORMULATION

From Professor Gerd Infanger’s lecture on *Large-Scale Portfolio Optimization*, we formulated our factor model accordingly:

$$R = UF + \epsilon$$

$$U = [T \times h] \text{ factors, } F = [h \times n] \text{ factor loadings} \\ \epsilon = [T \times n] \text{ residual returns}$$

$$\hat{Q}_H = \frac{1}{T} R^T R = \frac{1}{T} (F^T U^T U F + \epsilon^T \epsilon + F^T U^T \epsilon + \epsilon U F)$$

$$\hat{Q}_{HF} = F^T \hat{Q}_H F + D, \quad D = \text{diag}(\sigma_i)^2$$

Where \hat{Q}_{HF} represents the estimated covariance matrix of the factor model. The diagonal matrix D serves as the idiosyncratic risk for each individual security while the $F^T \hat{Q}_H F$ conjunction accounts for the systemic risk of the securities. R represents the factor returns as prescribed by the factor loadings, factor values, and residual returns. To find R , we found U and F first. In order to explore the possibilities for our U (and resultant F) matrix, which would be most predictive for our factor model, we decided to obtain an initial set of 18 factors.

The data we needed was pulled from the Federal Reserve Economic Data (FRED) website for the decade between Jan. 2003 and Jan. 2013. We downloaded monthly interval data to match our monthly

return data, and we split the data into sets for training (Dec. 2003 - Nov 2008) and testing (Dec 2008 - Nov 2013). Then, we converted the data into monthly percent change data, making it ready for use in creating our matrix F of factor loadings for each security and in building our matrix U of factor values. We chose the following factors to be representative of several sectors of the U.S. economy such as transportation, finance and real estate among others:

- | | |
|--------------------------------|---------------------------------|
| 1. 1-Month LIBOR Rate | 10. WTI Crude Barrel Price |
| 2. 3-Month T-Bill Yield | 11. Brent Crude Barrel Price |
| 3. 5-Year T-Bill Yield | 12. All Forms. US Gas Price |
| 4. 30-Year Mortgage Fixed Rate | 13. Gold Fixing Price |
| 5. TED Spread | 14. Industrial Production Index |
| 6. Unemployment Rate | 15. UoM Consumer Sentiment |
| 7. Consumer Price Index | 16. VIX Monthly Average |
| 8. \$/£ Exchange Rate | 17. US Trade Balance |
| 9. Personal Savings Rate | 18. S&P500 Index |

Given these initial factors, we decided to apply factor reduction techniques to help whittle down the overall number of factors in the final model.

III. FACTOR ANALYSIS

The goal in reducing the number of factors was simply to balance the tradeoff between the model's goodness-of-fit and its complexity. We attempt to explain as much variance as possible within the data, while avoiding poor predictions that often stem from overfitting.

To achieve this, we chose to implement backward stepwise AIC regression, backward stepwise BIC regression, and a hand selection technique.

A. Backward Stepwise AIC Regression

For each of the 500 stocks within our universe, we built an initial multivariate linear regression full model, complete with all 18 factors, on the training set of our data. Next, we calculated the Akaike Information Criterion (AIC) for the full model, where $AIC = 2k - 2\ln(L)$. Here, k represents the number of factors and L represents the maximum of the model's likelihood function. While AIC does not give us an idea of how good a model is in the absolute sense, it can demonstrate the value of a model in comparison to other models. Therefore, to find the best model, we systematically eliminated factors from the model, compared the relative AIC values, and chose the final model with the minimum AIC. All permutations of the 18 factors were given consideration in this backward stepwise AIC process.

Since we repeated this process for each of the 500 stocks, we ended up with 500 different models of stock behavior, each with a varying number of factors in the final model. We placed the model coefficients (intercept included) into a 500×19 matrix M_{AIC} and set the $M_{AIC(i,j)} = 1$ if a factor j remained in the final model for stock i , and $M_{AIC(i,j)} = 0$ if a factor j were eliminated from the final model of stock i . We ordered the resulting column-wise

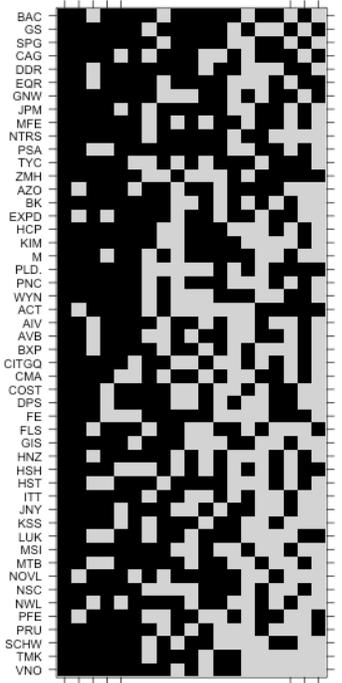


Fig. 1 - Visualization of M_{AIC}

sums from greatest to least to extract which factors

most frequently remained in the final model, and which factors were most frequently eliminated (*Table 1*). The visualization in *Fig. 1* represents 50 randomly chosen stocks from M_{AIC} using black squares to show which factors stayed, and gray squares to show which were eliminated, thereby giving a quick demonstration of the degree of AIC selectivity.

Paring down factors using AIC as a criterion did not result in a particularly strong reduction, seeing as 3 very highly correlated factors, each representing oil prices, remained within the top 5 most frequently occurring macro factors. Therefore, to build the full final model from this AIC reduction, we chose the top 8 most frequent factors. In order of highest to lowest appearance frequency, these were U.S. Gas Price Per Gallon, VIX Monthly Avg., Personal Savings Rate, WTI Crude Oil Price per Barrel, Brent EU Crude Oil Price per Barrel, 3 Month T-Bill Yield, TED Spread, and S&P 500 Index (*Table 1*). These factors gave us a fuller macroeconomic picture than simply the top 5, and helped improve the model's predictive ability. The specific characteristics of this model selection will be discussed in further detail in Section IV. In summary:

$$AIC_R^1 = \beta_0 + \beta_1 (GPG) + \beta_2 (VIX) + \beta_3 (PSR) + \beta_4 (WTI) + \beta_5 (BEU) + \beta_6 (3MY) + \beta_7 (TED) + \beta_8 (S\&P)$$

¹ $AIC_R = \beta_0 + \beta_1$ (U.S. Gas Price Per Gallon) + β_2 (VIX Monthly Avg.) + β_3 (Personal Savings Rate) + β_4 (WTI Crude Oil Price per Barrel) + β_5 (Brent EU Crude Oil Price per Barrel) + β_6 (3 Month T-Bill

B. Backward Stepwise BIC Regression

In addition to using AIC as a model selection criterion, we decided to test the Bayesian Information Criterion (BIC) too, where $BIC = \ln(n) * k - 2\ln(L)$. Again, k represents the number of factors and L represents the maximum of the model's likelihood function with n representing the number of observations. This criterion aims to perform the same function as AIC. Namely, it provides a value with which we can compare models (though it still does not give an absolute sense of the model's value). Again, in this selection process, we begin with a full model of 18 factors and iterate through all permutations of the factors in an attempt to find the model with the minimum BIC. The sole difference between AIC and BIC is a penalty of 2 vs. a penalty of $\ln(n)$ on the number of factors, k . With $n = 60$, as in our case, this penalty became twice as strong for BIC as AIC. The goal of using this alternate method was to see if the increased penalty on number of parameters would result in significant changes to the factors in our final BIC model as

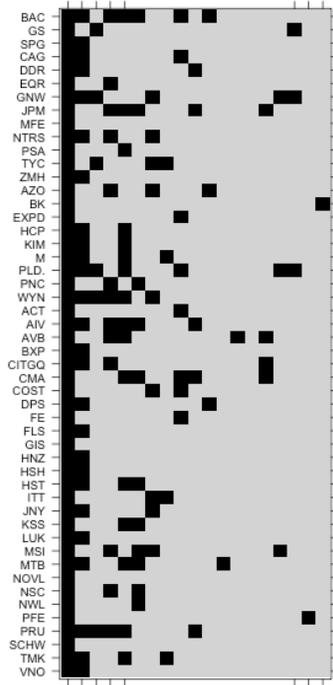


Fig. 2 - Visualization of M_{BIC}

compared to our AIC model. If so, we wanted to track these changes and their effects on predictive power and portfolio optimization.

We ran the backward stepwise BIC regression on each of the 500 stocks and calculated the matrix M_{BIC} in the same manner as M_{AIC} from before. The resulting M_{BIC} subset visualization of the same 50 stocks used to create the graphic of the M_{AIC} subset can be found in Fig. 2. From the figure, it is clear that as a direct result of the stronger penalty, the BIC stepwise regression concluded with far fewer overall occurrences of factors in its final models. In fact, not only did using BIC result in lower factor frequencies, but also fewer redundant factor appearances (Table 1).

Rather than being dominated at the top by several highly correlated measures of oil price, the BIC method chose a wider variety of factors.

AIC FREQUENCIES

1	X.Intercept.	500
2	US.Gas.Price...Gallon	260
3	VIX.Monthly.Avg	247
4	Personal.Savings.rate	246
5	WTI.Crude.Oil.Price.Barrel	230
6	Crude.Oil.Brent.EU...Barrel	226
7	Yield.3mo.T.Bill	222
8	TED.Spread	208
9	SPIndex	171
10	Yield.5yr.T.Bill	169
11	Gold.Fixing.Price	163
12	US.Trade.Balance	147
13	CPI	144
14	UoM.Consumer.Sentiment	134
15	US.UK.FX.Rate	129
16	Unemployment.Rate	127
17	Industrial.Production.Index	125
18	X1mo.LIBOR.Rate	125
19	Mortgage.Rate.30yr.Fixed	112

BIC FREQUENCIES

1	X.Intercept	500
2	Industrial.Production.Index	118
3	WTI.Crude.Oil.Price.Barrel	113
4	Personal.Savings.rate	91
5	SPIndex	87
6	TED.Spread	63
7	UoM.Consumer.Sentiment	48
8	US.Gas.Price...Gallon	45
9	Unemployment.Rate	38
10	US.UK.FX.Rate	35
11	CPI	31
12	Crude.Oil.Brent.EU...Barrel	31
13	Gold.Fixing.Price	31
14	Yield.3mo.T.Bill	28
15	X1mo.LIBOR.Rate	25
16	VIX.Monthly.Avg	17
17	US.Trade.Balance	15
18	Mortgage.Rate.30yr.Fixed	10
19	Yield.5yr.T.Bill	7

Table 1 - Factor Frequencies of Stepwise AIC & BIC Regressions

Therefore, we deemed the top 5 factors given by the automated BIC method to be sufficient and representative indicators of the economy for our final model. In order of highest to lowest frequency, these factors were Industrial Production Index, WTI Crude Oil Price per Barrel, Personal Savings Rate, S&P 500 Index and TED Spread. The specific characteristics of the model will be discussed in Section IV. In summary:

$$BIC_R^2 = \beta_0 + \beta_1 (IPI) + \beta_2 (WTI) + \beta_3 (PSR) + \beta_4 (S\&P) + \beta_5 (TED)$$

C. Hand Selected Factors (H-S)

The last method we used to choose factors from the original 18 was a simple hand selection method. In order for multivariate linear regression to be an effective choice for model building, one major assumption is that the factors we chose as independent variables are, in fact, independent. When looking at the macroeconomic data we collected, it is clear that this assumption does not hold true in the strict sense of independence (i.e. 0 correlation). However, this requirement of independence was approximately true in some cases. By visualizing the factor correlation matrix, found in Fig. 3, we found that a reasonable range to approximate factor independence would be

² β_5 (Brent EU Crude Oil Price per Barrel) + β_6 (3 Month T-Bill Yield) + β_7 (TED Spread) + β_8 (S&P 500 Index)

² $BIC_R = \beta_0 + \beta_1$ (Industrial Production Index) + β_2 (WTI Crude Oil Price per Barrel) + β_3 (Personal Savings Rate) + β_4 (S&P 500 Index) + β_5 (TED Spread)

to select factors whose correlations fell roughly within the $[-0.2, 0.2]$ range.

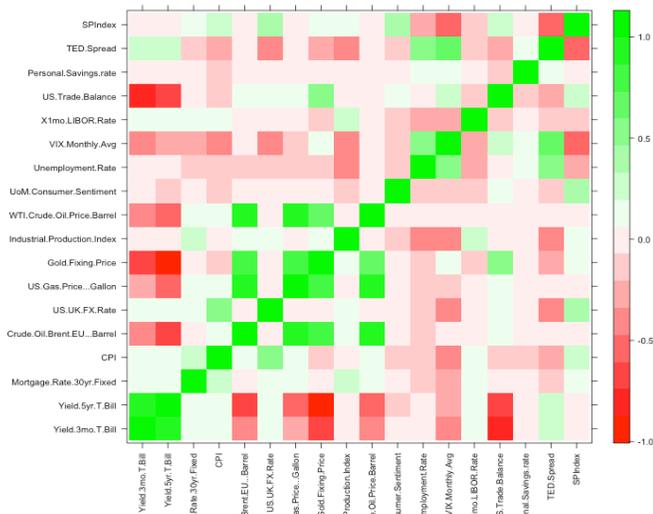


Fig. 3 - Correlation Matrix of Macroeconomic Factors

Now, in addition to examining the mathematical correlations between factors, we attempted to choose factors in such a manner that would help comprehensively explain stock behavior in multiple sectors of the economy. We posited that the behavior of the S&P 500 should be incorporated, as should economic output across multiple sectors, price of consumer goods, price of important commodities and the current state of the U.S. Treasury. By incorporating these major economic facets, we aimed to capture as much variance in various stocks as possible, while keeping the model simple and flexible enough to adapt to new data.

Putting together the hypotheses of both the correlations between factors and economic behavior, we decided upon a final 5 macroeconomic factors for the hand selected model. Namely, we chose the S&P 500 Index, the Industrial Production Index, the Consumer Production Index, the WTI Crude Oil Price per Barrel and the 3-Month Treasury Yield.

$$H-SR^3 = \beta_0 + \beta_1 (S\&P) + \beta_2 (IPI) + \beta_3 (CPI) + \beta_4 (WTI) + \beta_5 (3MY)$$

³ $H-SR = \beta_0 + \beta_1 (S\&P\ 500\ Index) + \beta_2 (Industrial\ Production\ Index) + \beta_3 (Consumer\ Price\ Index) + \beta_4 (WTI\ Crude\ Oil\ Price\ per\ Barrel) + \beta_5 (3\ Month\ T-Bill\ Yield)$

IV. REGRESSION ANALYSIS

After finding these three sets of parameters, we had 3 distinct models, to which we applied to our return data. Therefore, as a baseline set of models, we built each model on the 60 months of training data (Dec 2003 - Nov 2008) and predicted on the 60 months of testing data (Dec 2008 - Nov 2013) to gauge each model's fit. In the end, we ran each model 500 times and found some summary statistics to help compare the performance of each model.

A. Baseline Model Summary Statistics

Below, Table 2 presents a few of the results we found important to the model's fit. In order of the table, we calculated the mean of correlations between the training set and the model's residual returns (Training Mean Cor.), the mean of correlations between the testing set and the model's predicted returns (Predicted Mean Cor.), the percentage of model predictions which were positively correlated with the training set above the .05 threshold (Predicted Cor. > 0.05), the overall mean of each model's predicted-vs.-actual root mean squared error (Pred. Mean RMSE) and the mean of the R^2 values for the models.

	Training Mean Cor.	Predicted Mean Cor.	Predicted Cor. > .05	Pred. Mean RMSE	Mean R^2
AI C	0.500	0.051	52.6%	14.47	0.26
BIC	0.444	0.096	62.6%	9.79	0.21
H-S	0.410	0.168	74.0%	10.12	0.18

Table 2 - Statistics from the AIC/BIC/H-S Regression Models

We looked at graphs of modeled returns vs. actual returns (Fig. 4 & Fig. 5) and correlations in order to understand how successful the model was at predicting stock returns. If we found a huge number of models with negative or zero correlations, we would have reason to worry. However, we see a majority of models with positive correlations and, in fact, a majority of models with correlations above .05. This percentage rises significantly from AIC to BIC to the H-S model, along with our predicted mean correlation. We also see the expected drop-off between in-sample and out-of-sample mean correlations. Working to maximize this amount while minimizing the difference could certainly help improve our model.

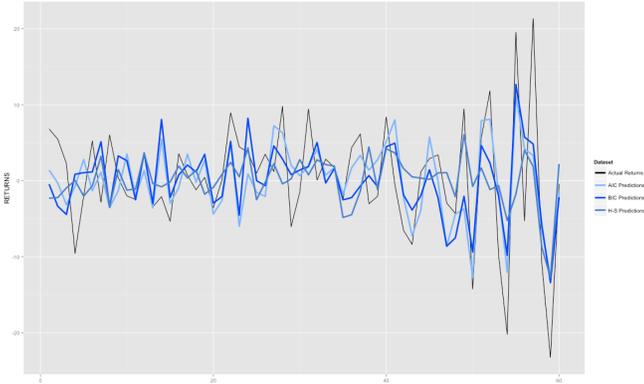


Fig. 4 - JPM Training Set Returns (black) vs. Model Returns (blue)

Similarly, we see the stronger correlations in predictions matching up with a decrease in predicted mean RMSE, albeit with a slight up-tick for the H-S model. Lastly, we examine the R^2 values for each model. These values are relatively low and the number of factors in play certainly has an effect on the higher R^2 value for the AIC model (8 factors vs. 5 for BIC and H-S). The weak values of the AIC model's actual prediction performance lend credence to the hypothesis that its inflated R^2 value is a result of its increased number of factors rather than the model having better, more predictive factors.

Overall, the H-S model produced the strongest predicted correlations, although it could benefit from additional factors, as indicated by its lower R^2 . The BIC performed decently, though its factors turned out to be weaker predictors (particularly Personal Savings Rate and TED Spread - the only 2 differing factors from H-S model). Lastly, the AIC suffered from a combination of over-fitting to the training set and having weaker, redundant factors.

In terms of the portfolio optimization formulation, we took the respective U 's provided by these models and multiplied by F (training period data) to find the matrix R_F of returns explained by our factors. We calculated \hat{F} by simply predicting the model one time period ahead. Finally, we calculated D by taking $R - R_F$ then calculating the standard deviation of each of the 500 columns, placing those 500 values in the diagonal of 500x500 square matrix and squaring the matrix. From this analysis we therefore obtained all parts of the matrix necessary to run our long run portfolio optimization.

B. Rolling Window Optimization

After this analysis of the baseline set of models, we wanted to tune a more precise optimization calibration. To achieve this, we determined that instead

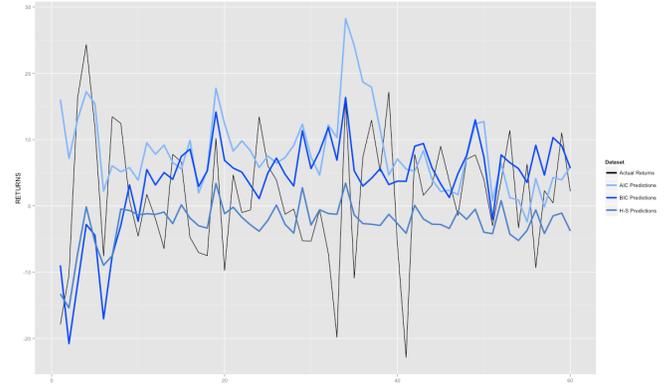


Fig. 5 - JPM Test Set Returns (black) vs. Predicted Returns (blue)

of a long run optimization, we could shift funds in the short term to obtain a better strategy. Mainly, we wanted to have the ability to make slight adjustments to our portfolio in response to the S&P 500's behavior. We chose, therefore, to build our matrix of returns R_F on a rolling 60-month window of iteratively demeaned data, wherein at each time period we would calculate a new U and shift F appropriately to include the correct months. From the new R_F we calculated a similarly new vector \hat{F} and matrix D . We performed 60 iterations, moving ahead by one month each time to go from a model built on training data from Dec 2003 - Nov 2008 to a final model built on training data from Dec 2008 - Nov 2013. At each time period we ran a new portfolio optimization to see the effect of the new data on our portfolio allocation. We did this same rolling window analysis for each of the AIC, BIC and H-S models.

While we did not perform as extensive statistical analysis on each of these 60 rolling window models as on the baseline models, we continued to track cor-

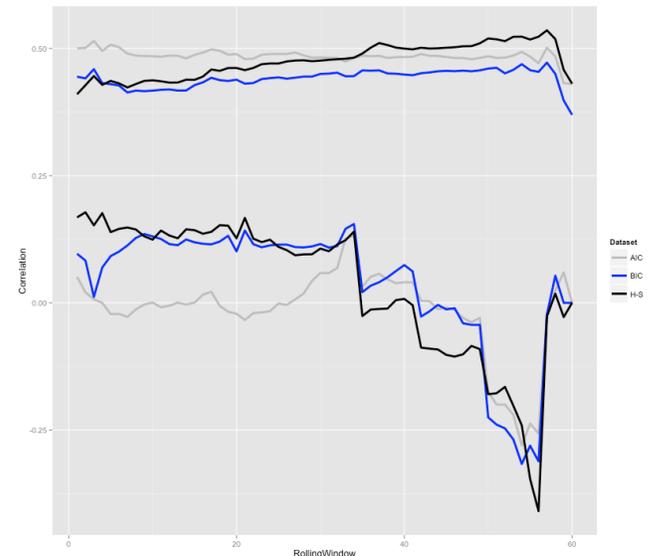


Fig. 6 - Rolling Window In-Sample & Out-of-Sample Correlations

relations for both the in-sample data and the out-of-sample data. The data for the in-sample and out-of-sample correlations are shown in *Fig. 6*. In-sample data are clustered at the top, while out-of-sample correlations hover nearer to zero. The gray, blue and black lines correspond to AIC, BIC, and H-S models.

Unsurprisingly, the mean correlation between our predicted values and actual testing return data starts to lose quality after the 35th model and even more so after about the 50th model. This is mainly due to the fact that our test dataset has become so small that variation easily overwhelms any sense of the mean. Therefore, our models do not track well with the constantly reducing test sets. The in-sample correlations, on the other hand, continue to fall within similar ranges as we change the range of the sample, indicating to some degree that our models deal fairly well with the new data and do not have much in-sample bias.

Given the relatively consistent performance of our AIC, BIC, and H-S models, we decided to go ahead with testing our short run portfolio optimization, the results of which are discussed in the next section.

V. PORTFOLIO OPTIMIZATION

To test the performance of each of the three models we set up an optimization problem in GAMS. Our optimization problem was setup to minimize the portfolio variance given a desired return and was of the form:

$$\begin{aligned} \text{Minimize} \quad & \frac{1}{2} (\mathbf{x}^T \mathbf{D} \mathbf{x} + \frac{1}{T} \mathbf{v}^T \mathbf{v}) \\ \text{s.t.} \quad & \mathbf{e}^T \mathbf{x} = 1 \\ & \hat{\mathbf{r}} \mathbf{x} \geq \boldsymbol{\rho} \\ & \mathbf{R}_F \mathbf{x} = \mathbf{v} \end{aligned}$$

Where \mathbf{x} is the optimal portfolio allocations to each stock, $\boldsymbol{\rho}$ is the desired return, and \mathbf{D} , \mathbf{R}_F , $\hat{\mathbf{r}}$ are as previously defined for each model.

We then ran each model through the optimization problem for all 60 months over a range of desired returns to be able to create an efficient frontier and see how well each of them performs. The resulting efficient frontiers can be seen in *Figure 7 - Efficient Frontiers*.

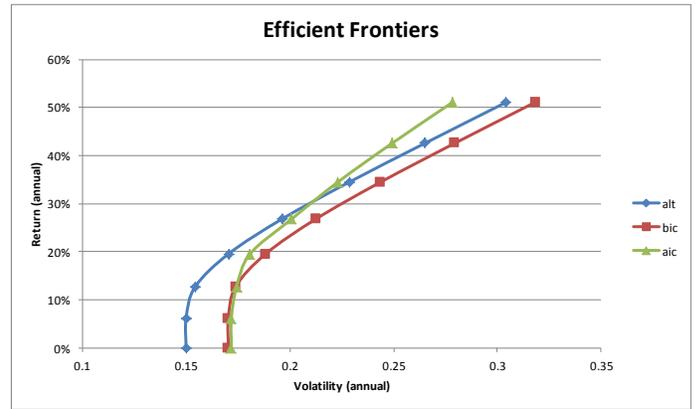


Fig. 7 – Efficient Frontiers

From these efficient frontiers we can see that our H-S model (alt) is better all-around than the BIC model because for each desired return it is supposed to get that return for less variance. The AIC model starts out with an efficient frontier similar to the BIC model, but for high desired returns seems to show relatively the lowest variance. Looking in to why/how the AIC model’s curve is so different from BIC and H-S models’ curves is something we have yet to do but it does warrant looking into.

After looking into the efficient frontiers and how our optimization problem says we should allocate our funds for each of the different models we wanted to see how those allocations actually performed, and since we have the returns information for each period we predicted over we were able to see what the realized returns would have been if we were to have followed our models and invested according to them. The results from checking what our realized returns would have been can be seen in *Figure 8*, and those as a percent of the target returns can be seen in *Figure 9*.

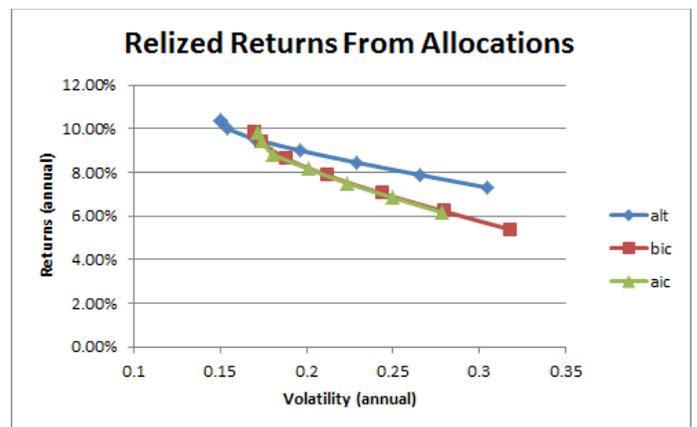


Fig. 8 – Realized Returns from Optimization Allocation

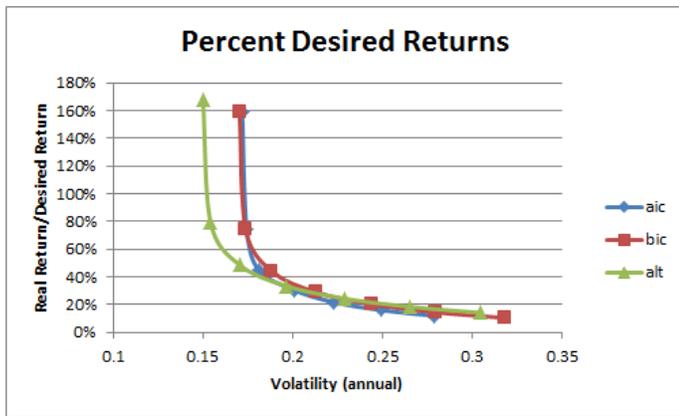


Fig. 9 – Realized Returns as a Percent of Desired Returns

We seem to get fairly good returns when we set our targeted returns to be fairly low, and as we increase our target returns we get lower and lower realized returns, which is exemplified best in Figure 9, where as we increase our target returns we get exponentially a smaller and smaller percent of that target. This is troubling to see since you would hope for the opposite, increasing realized returns when you set your target returns to be higher. In continuing studies on our models trying to figure out what is causing this would be very important and help building better and more robust models in the future.

VI. STATISTICAL FACTOR MODEL

We also compared our H-S model against the statistical factor model that was included with our data. We only had data to predict the statistical factor model for the first month of our test data, so we optimized both the Statistical Factor Model and the H-S model just for the first month, and compared their efficient frontiers. Our H-S model does provide lower variance at lower returns, but quickly gets outperformed by the statistical factor model, as target returns increase.

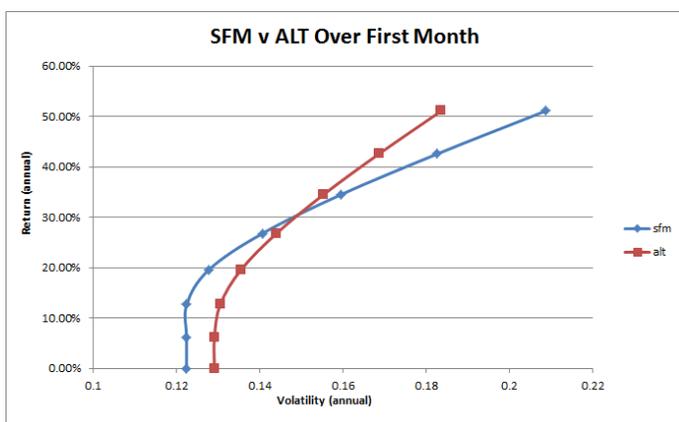


Fig. 10 – SFM v H-S Model One Month Efficient Frontier

VII. CONCLUSION

This study provided us insight on multiple fronts.

Our regressions were able to seek out the most relevant factors to the stocks, possibly providing predictive power of the macroeconomic method. However, despite these methods, we had two main curiosities in our study. First, the hand-selected method proved more predictive than either the AIC or BIC methods. Second, our realized returns from our optimized allocation varied inversely with volatility. Further research is needed to understand why each of these unexpected results occurred. Finding more strongly predictive factors is likely the first step in addressing these issues.

As for the second step, we believe that the model may have a hard time adjusting to volatility in the market. Part of this suspicion comes from the model's abysmal performance in March 2009, when the 2008 recession bottomed out on the stock market. The macroeconomic model reported no better than a 12% loss during this month. That said, further work in adjusting the time window for the factor model would be interesting. It is possible that the five-year window was too long for a factor model as the United States' real estate and financial landscapes changed radically in the post-recession period. In essence, when something becomes the "new normal," a five-year window will have a hard time adjusting to that normal.

However, adjusting the window is one of many adjustments that could be made. Altering the leverage ratio to allow an increased capital supply to invest would be interesting as well as other minor and major parameter tweaks. Unfortunately, we did not have the resources to do a sensitivity analysis on parameters like leverage ratio, time window, and single-security maximum stakes.

Additionally, mixing up the security mix from the United States' S&P 500 to include a more international mix of stocks as well as potentially corporate and municipal bonds could help diversify away the idiosyncratic systemic risk of the S&P 500 itself.

Lastly, there are many more than eighteen macroeconomic factors to choose from, and a selection and analysis could always prove more interesting.

While we were disappointed to not be able to outperform the statistical factor model (as Connor predicted), we thoroughly enjoyed the experience of working within the macroeconomic environment and

hope that our analysis regarding factor selection and stochastic reformulation can aid future studies.

VIII. ACKNOWLEDGEMENT

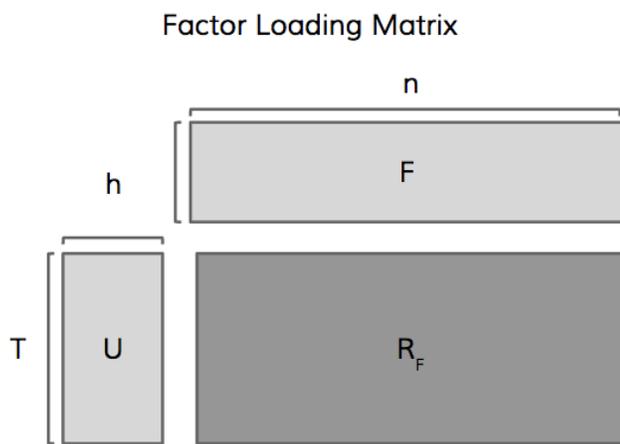
We would like to thank Professor Gerd Infanger for his guidance in this work. We would also like to thank the people at GAMS for helping us through multiple coding difficulties.

IX. APPENDIX

Figure A.

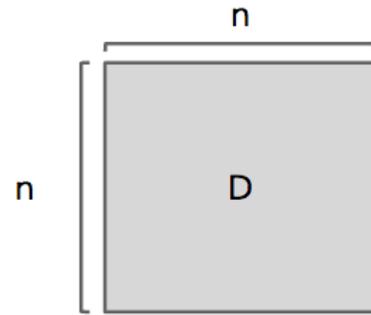
Factor Model Type	Inputs	Estimation Technique	Outputs
Macroeconomic	<i>Security Returns and Macroeconomic variables</i>	<i>Time-series regression</i>	<i>Security factor betas</i>
Statistical	<i>Security returns</i>	<i>Iterated time-series/ cross-sectional regression</i>	<i>Statistical factors and security factor betas</i>
Fundamental	<i>Security returns and security characteristics</i>	<i>Cross-sectional regression</i>	<i>Fundamental factors</i>

Figure B.

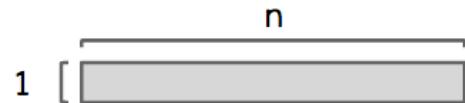


Where: T = number of time periods
 h = number of factors
 n = number of stocks

Diagonal Variance Matrix



Expected Returns (\hat{r})



X. REFERENCES

- Connor, Gregory. "The Three Types of Factor Models: A Comparison of Their Explanatory Power." *Financial Analysts Journal* 51.3 (1995): 42-46. *Stanford Coursework*. Web. 19 Feb. 2015.
- Infanger, Gerd. "Large-Scale Portfolio Optimization." MS&E 348 Class. Stanford, California. 8 Jan. 2015.