

# Constructing an ETF Portfolio

IAN L. KAPLAN

[iank@bearcave.com](mailto:iank@bearcave.com)

<http://www.bearcave.com>

December 2, 2014

## **Abstract**

These notes discuss the construction of investment portfolios consisting of Exchange Traded Funds, ETFs. These portfolios are constructed using quantitative portfolio techniques. This document is somewhere between an academic paper and a lab notebook. As a result, this paper is written in a more informal style. As a record of experimentation, these notes are a "warts and all" account. They cover the process of developing the portfolio, including back-testing mistakes and the experiments that didn't work out.

Nothing in these notes should be viewed as investment advice.

## Introduction

These notes discuss the construction of an investment portfolios constructed from Exchange Traded Funds (ETF).

Why ETFs and not other assets?

ETFs are constructed with a wide variety of market exposures. For example, there are ETFs that attempt to replicate indexes like the S&P 500 and the S&P Mid-cap 400 indexes. There are ETFs that provide negative market exposure and leveraged long exposure. There are also specialized ETFs in a wide variety of sectors (software, semiconductors, market indexes, precious metals).

The number of ETFs has exploded since they were first introduced in 1993. One drawback of ETFs is that many of them, especially the newer exotic ETFs, have a short history (sometimes only a year or less).

Unlike index mutual funds, ETFs trade like stocks. A portfolio constructed from a few ETFs have some level of diversification, since an ETF itself is composed on a portfolio of assets. ETFs also have tax advantages compared to mutual funds or index funds (unlike mutual funds, trading within an ETF does not incur tax liability).

## Past is Prologue

The approach used here in constructing investment portfolios assumes that the past has some influence on the future. This includes the assumption that there is serial dependence in asset returns and that assets that have increased in value in the recent past have a higher likelihood of doing so in the next quarter.

Empirically, the returns in the universe of ETF assets are not normally distributed. This is especially true for the back-test period used here, which includes the 2008 crash where a number of assets have (negative) returns that would be far in the tail of a normal distribution.

## **Challenges with ETFs**

### **Liquidity**

Some ETFs are thinly capitalized with only a few million dollars worth of net assets. These ETFs also tend to be thinly traded. For example, the net assets of the LSTK ETF (iPath Pure Beta Livestock ETN) were only 2.09 million in November 2014 and the three month average volume is only 788 shares. The LSTK ETF had a year-to-date return in November of 2014 of 26.88 percent. However, the low liquidity and the potential market impact in trading a low liquidity ETF like LSTK makes this a riskier investment than the return statistics suggest. The etf.com screen used to select the current ETF universe filters out ETFs like LSTK.

### **Exotic Market Behavior**

ETFs package a wide variety of market exposure. This includes currency ETFs like YCS (ProShares UltraShort Yen).

For the three months ending on November 28, 2014, YCS had a 29.4 percent return. In the same three month period ending on November 28, 2014 the exchange rate between the dollar and the Japanese yen went from about 104 yen/dollar to about 118 yen/dollar, causing YCS to increase proportionally in value. The yen/dollar exchange rate was at a ten year high (for the dollar and a low for the yen) of about 120 yen/dollar.

For the appreciation in YCS to continue, the exchange rate would have to rise (relative to the dollar) to levels that have not been seen since 1998. The 1998 exchange rate was also accompanied by extreme volatility. In 1998 the yen/dollar exchange rate fell about 34 percent in two months, from July to August. These factors make YCS an inappropriate choice for the type of portfolio that is being constructed in this paper.

## Holding Period

A quarterly holding period is used for the portfolios. This time period allows the portfolio to be adjusted for new market conditions.

To provide sufficient data for back-testing, weekly return data is used. Weekly returns do not have exactly the same distribution as monthly or quarterly returns, so this is an imperfect solution.

## Calculating Returns

In academic finance, continuously compounded log returns are commonly used:

$$r_t = \log(P_t) - \log(P_{t-1})$$

Calculating returns in this manner is correct if returns have a distribution that is close to normal. This is not the case for the ETF returns during the back-test period which includes the market crash of 2008. For example, SKF (ProShares UltraShort Financials) dropped from 753.00 to 190.08 over a quarter.

$$r_t = \log(P_t) - \log(P_{t-1}) = \log(190.08) - \log(753.00) = -1.37662$$

This is a loss greater than the value of the asset, which is not possible with stocks (or ETFs).

A better way to estimate returns for non-normal data is to use arithmetic returns:

$$R_t = \frac{P_t}{P_{t-1}} - 1$$

Using arithmetic returns  $R_t = -0.7475697$ , which accurately estimates the loss.

This approach to estimating returns is also a better model for how an investor behaves, compared to compound returns. An investor buys the assets at the start of the period and then sells at the end of the period.

Arithmetic returns vs. compound returns are discussed in the paper *Linear vs. Compounded Returns Common Pitfalls in Portfolio Management* by Attilio Meucci, available on SSRN.

## Back-testing

Some people emphatically proclaim that they never back-test. They point out that back-test results are never the same as out-of-sample actual results. So why bother?

Perhaps the core reason to back-test is to catch problems in the construction of the portfolio construction algorithm. For example, in the gain-loss portfolio below, assets are pre-filtered by the gain-loss ratio. Without back-testing I might not have noticed that assets with similar gain-loss ratios tend to be similar ETFs (for example, Oil or Energy ETFs). To avoid a portfolio heavy in a single asset type (energy) or sector, I added code to filter out assets with high paired correlation.

As the Mutual Fund adds say, past performance is not indication of future performance. But past performance does provide an example of how the algorithm behaves. If the algorithm performs poorly in the back-test it is unlikely to suddenly perform much better with out-of-sample data.

There is a lurking danger in back-testing, however. There is rarely enough financial data for both back-testing and out-of-sample tests. This is especially true of ETFs, which are relatively new investment instruments.

The back-tests here use a period of about seven years. This means that there is no out-of-sample data except for the future.

When back-testing a portfolio construction algorithm, features of the algorithm are adjusted to yield better performance. In doing this, future information is applied to the past. If this is done frequently enough the result will be an algorithm that performs well on past data, but doesn't necessarily do well on the out-of-sample future data. This is over-fitting.

A thoughtful discussion of the pitfalls of back-testing can be found in *Pseudo-Mathematics and Financial Charlatanism: The Effects of Back-test Over-fitting on Out-of-Sample Performance* by David H. Bailey, Jonathan M. Borwein, Marcos Lopez de Prado, and Qiji Jim Zhu available on SSRN. This paper also proposes a way to estimate some of the inaccuracy introduced by back-testing.

## Portfolio Optimization

Portfolio optimization involves finding the asset weights that yield highest portfolio return for a given level of risk. Calculating an optimal portfolio for a particular level of risk requires an estimate of the risk and forecasted (future) return.

### Risk

Risk is a retrospective measure that is reported as variance, value-at-risk (VaR) or Conditional Value at Risk (CVaR, also known as Expected Shortfall or Extended Tail Loss). One problem with risk estimates is that they are difficult to verify in back testing where the actual "future" is known.

CVaR is, theoretically, a more attractive risk measure since it measures down-side risk, without penalizing upside volatility. A problem with CVaR is that it relies on an estimate of the negative tail of the distribution where, empirically, there are few values. To address this problem a normal distribution is sometimes used. The back-test period includes a market crash that contains extreme values. A normal distribution differs dramatically from this empirical distribution. Another solution is to fit a distribution (perhaps using R's `density` function) and use the tail of this distribution to estimate the CVaR value.

Each portfolio is calculated from a year of past weekly data. With only a year of weekly data the error in CVaR estimate is so high that any advantage it has over using variance ( $\sigma^2$ ) is lost.

## Forecasting Future Returns

In classic (Markowitz) mean-variance portfolio optimization the future return is forecast by the mean return over a historical window. This assumes that the mean is stable (which is only the case over a very long or very short period).

To forecast future returns, one year of weekly past data (e.g., 52 values) is used to forecast the return twelve weeks ahead.

The portfolio is bought at the "current week" based on the portfolio estimate using the past year of weekly data. The portfolio is held for 12 weeks and then sold realizing a return.

Some functions can be used to forecast future return are:

- mean
- exponential moving average (EMA) - from R's TRR package
- linear mean (e.g., the value predicted by an ordinary least squares line through the past returns).

## Portfolio Weights

In the paper, long-only portfolios are constructed. The analytic mean-variance equations cannot be used (since mean-variance may short sell assets). One way to calculate an optimal long-only portfolio is to use R's `solve.QP` in the `quadprog` package <sup>1</sup>.

The `solve.QP` function minimizes:

$$\frac{1}{2} \mathbf{b}^T \mathbf{D} \mathbf{b} - \mathbf{b}^T \mathbf{d} \text{ such that } \mathbf{A}^T \mathbf{b} \geq \mathbf{b}_0$$

If there are  $n$  stocks with  $T = 52$  time periods

- $\mathbf{D} = 2\Sigma = 2 \cdot \text{cov}(\mathbf{R})$  an  $n \times n$  matrix

---

<sup>1</sup>The material in this section is based on Guy Yollen's lecture notes for *Financial Data Modeling and Analysis in R*, University of Washington, 2012

- $\mathbf{b} = \mathbf{w}$  (the portfolio weights) a vector of length  $n$
- $\mathbf{d} = 0$
- $\mathbf{A}$  is a constraint matrix ( $n \times m$ )
- $\mathbf{b}_0$  is a vector of constraint bounds (length  $m$ )

There are  $n$  stocks with a forecasted return  $\mathbf{r} = r_1, \dots, r_n$ . The portfolio return is  $r_p = \mathbf{w}^T \mathbf{r}$

For a portfolio where

- weights sum to 1
- portfolio return target is  $r_p$
- no short sales (e.g., all portfolio weights are greater than 1)

$$\mathbf{A} = \begin{bmatrix} 1 & r_1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & r_2 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 1 & r_n & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\mathbf{b} = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \end{bmatrix}$$

$$\mathbf{b}_0 = \begin{bmatrix} 1 \\ r_p \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

The  $\mathbf{b}_0$  vector contains the portfolio target return,  $r_p$ .

For each portfolio, a target return must be chosen. If this return can be realized, the optimizer will choose portfolio weights that realize this return, with the lowest risk (variance).



In order to choose a maximum return that is likely to be realized the equation below is used for the target return:

$$r_{target} = max(\mathbf{r}_p) - max(\mathbf{r}_p) \times frac$$

Where  $\mathbf{r}_p$  is a vector of forecasted returns for the portfolio assets and  $frac$  is usually between 0.1 and 0.3.

An alternative picking a single target return is to calculate an efficient frontier for the portfolio by constructing portfolios for a range of values. The portfolio with the maximum Sharpe ratio is used in some of the models discussed below.

One of the challenges of using `solve.QP` is getting the arguments right. There is a GNU port of an optimization language that can be used to describe the constraints that is much clearer than the vector definitions used for `solve.QP`. The R-blogger Systematic Investor has written a post that describes this:

Portfolio optimization constraints with the GNU MathProg language

## A Hand Selected ETF Portfolio

The first portfolio discussed consists of ETFs that were selected "by hand". Selecting ETFs for this portfolio was an iterative and time consuming process. The portfolios discussed later in this report use software screening to select a set of ETFs from a large ETF universe.

In an early version of this portfolio, ETFs with negative exposure were included. For example, ETFs with -1x S&P 500 exposure and -1x oil and gas exposure. These negative  $\beta$  assets were selected because these assets are expected to have positive return when the market went down. The speculation was that ETFs with negative beta would be similar to a short position, without the complexity of a short position.

The portfolios calculated here are all long-only portfolios. Unlike long-short portfolios, a long-only portfolio does not profit from a market drop. It was

hoped that negative beta ETFs would act like a short position and profit from market down turns. However, the portfolio optimizer never chose these assets, even when the market went down.

The negative beta ETFs attempt to replicate a negative market position on a daily basis. Holding a negative beta S&P 500 ETF for month is not the same as holding a short position in the S&P 500 for a month because the ETF is re-balanced.

The ETFs that were selected by hand for the first portfolio are listed below in Table 1.

ETF Symbol	Description	ETF Type	Inception Date
SHY	Core short-term US Bond	Fixed Income	2002-07-22
AGG	Core Total US Bond Market	Fixed Income	2003-09-22
HYG	High Yield Corporate Bond ETF	Fixed Income	2007-04-04
SPY	S&P 500	Equity	2000-07-24
XLU	Utilities Select Sector	Equity	1998-12-16
XPH	S&P Pharmaceuticals	Equity	2006-06-19
XSD	S&P Semiconductor	Equity	2006-01-31
XRT	S&P Retail	Equity	2006-06-19
XHB	S&P Homebuilders	Equity	2006-01-31
XOP	Oil and Gas	Equity	2006-06-19

Table 1: ETF Universe

In a portfolio of ETFs, the back-test time period is set by the ETF with the most recent inception data. In this portfolio this is the HYG High Yield Corporate Bond ETF. The inception date is 2007-04-04.

The correlation plot for the portfolio ETFs is shown below in Figure 1




Figure 1: Correlation of Asset Returns

The portfolio  $\beta$  between the ETF funds and the S&P 500 ETF (SPY) is shown below in Table 2. The  $\beta$  is calculated from weekly returns.

ETF Sym	ETF Description	Beta w/S&P 500
SHY	Core short-term US Bond	-0.0241
AGG	Core Total US Bond Market	0.0461
HYG	High Yield Corporate Bond ETF	0.4828
XLU	Utilities Select Sector	0.6323
XPH	S&P Pharmaceuticals	0.8063
SPY	S&P 500	1.0000
XRT	S&P Retail	1.1252
XSD	S&P Semiconductor	1.1658
XOP	Oil and Gas	1.4292
XHB	S&P Homebuilders	1.5342

Table 2: ETF Beta with the SP 500

The weekly return distribution for the selected ETF assets is shown below:




Figure 2: ETF Weekly Return Distribution

Table 3 shows the weekly mean and standard deviation values (in percent).

ETF Sym	Description	Weekly Mean (percent)	StdDev (percent)
SHY	Core short-term US Bond	0.0442	0.1831
AGG	Core Total US Bond Market	0.0942	0.8219
XLU	Utilities Select Sector	0.1248	2.5119
HYG	High Yield Corporate Bond ETF	0.1327	1.8867
XHB	S&P Homebuilders	0.1411	5.5011
SPY	S&P 500	0.1542	2.8339
XSD	S&P Semiconductor	0.1824	4.2594
XOP	Oil and Gas	0.2318	5.0828
XRT	S&P Retail	0.2664	3.7614
XPH	S&P Pharmaceuticals	0.3440	2.8591

Table 3: Weekly ETF mean and StdDev (percent)

If the HYG high yield corporate bond ETF were dropped from the portfolio, there would be a larger back-test time period available (no high yield bond fund with an earlier inception data has been identified.) This ETF does have some attractive characteristics that make it difficult to discard. In particular it has good return and low volatility compared to the equity ETFs.

## Forecasting Returns

An important step in building a portfolio is forecasting the future returns that will be used in portfolio optimization. In classic mean-variance optimization the future return is estimated by the mean return.

Using a window of 52-weeks of past weekly returns, tables 4 and 5 show the correlation between the value forecasted by the function and the actual monthly and quarterly return.

	SHY	AGG	HYG	SPY	XLU	XPH	XSD	XRT	XHB	XOP
mean	0.2251	-0.0680	-0.0333	0.0876	0.0159	-0.0271	0.0272	0.0267	0.0338	-0.1173
EMA	0.0261	-0.0324	-0.1068	0.0838	0.0147	-0.1650	0.0109	-0.0300	-0.1241	-0.0088
ols	-0.3029	-0.1309	-0.0089	0.1132	0.1074	-0.0719	-0.0402	-0.0493	-0.1683	0.0269

Table 4: Correlation of Mean Functions with 4-week return

	SHY	AGG	HYG	SPY	XLU	XPH	XSD	XRT	XHB	XOP
mean	0.5057	-0.0223	0.0002	0.1261	0.1518	0.0607	0.1010	0.0690	0.0170	-0.1902
EMA	0.2742	-0.2764	-0.0033	0.3678	0.2586	-0.1758	0.3812	0.3043	-0.1861	0.0960
ols	-0.5387	-0.2958	0.1370	0.1593	0.0900	-0.1000	-0.0235	0.0383	-0.2537	-0.1313

Table 5: Correlation of Mean Functions with 12-week return

The correlation between the predicted return and the quarterly (12-week) return is much stronger for all functions than the correlation with the monthly (4-week) return. This suggests that a portfolio that trades quarterly may be a better choice than a portfolio that trades monthly.

Table 6 shows the function that produce the highest absolute correlation with quarterly future returns. Some of these correlations are negative correlations.

	function	correlation	beta sign
SHY	ols	negative	negative
AGG	ols	negative	postive
HYG	ols	postive	postive
SPY	EMA	postive	postive
XLU	EMA	postive	postive
XPH	EMA	negative	postive
XSD	EMA	postive	postive
XRT	EMA	postive	postive
XHB	ols	negative	postive
XOP	mean	negative	postive

Table 6: Functions vs. correlation and market beta

Selecting these functions and the associated adjustment factor is over-fitting, since information from the future is used in the past (when it could not have been known).

The mean return is negatively correlated with the future return.

As Table 6 shows, the bond funds SHY and AGG have a negative beta with the benchmark (which is to be expected with bonds). In the case where the beta is not negative, the negative correlation of the mean can be converted to a positive correlation by scaling the forecast by a multiplier:

$$\text{multiplier} = \text{sign}(\beta) \cdot \text{sign}(\rho) \quad \text{if } |\beta| \geq 0.1$$

where  $\rho$  is the correlation. The result is shown in Table 7

	function	multiplier
SHY	ols	1.00
AGG	ols	1.00
HYG	ols	1.00
SPY	EMA	1.00
XLU	EMA	1.00
XPH	EMA	-1.00
XSD	EMA	1.00
XRT	EMA	1.00
XHB	ols	-1.00
XOP	mean	-1.00

Table 7: Prediction Multiplier




Figure 3: Quarterly Portfolio Weights

The table below shows the numeric portfolio weights. Portfolio weights are rounded to two decimal places.

	SHY	AGG	HYG	SPY	XLU	XPH	XSD	XRT	XHB	XOP
2008-05-19	0.00	0.20	0.00	0.00	0.03	0.00	0.77	0.00	0.00	0.00
2008-08-11	0.20	0.00	0.00	0.00	0.00	0.00	0.54	0.26	0.00	0.00
2008-11-03	0.03	0.00	0.00	0.00	0.00	0.25	0.00	0.00	0.72	0.00
2009-01-26	0.00	0.25	0.00	0.00	0.00	0.00	0.00	0.00	0.75	0.00
2009-04-20	0.18	0.00	0.00	0.00	0.00	0.00	0.05	0.76	0.00	0.00
2009-07-13	0.00	0.00	0.52	0.00	0.00	0.00	0.48	0.00	0.00	0.00
2009-10-05	0.00	0.00	0.41	0.00	0.00	0.00	0.00	0.59	0.00	0.00
2009-12-28	0.00	0.00	0.38	0.00	0.00	0.00	0.62	0.00	0.00	0.00
2010-03-22	0.00	0.20	0.00	0.00	0.00	0.00	0.00	0.80	0.00	0.00
2010-06-14	0.00	0.04	0.00	0.00	0.93	0.00	0.00	0.00	0.03	0.00
2010-09-07	0.00	0.31	0.00	0.00	0.69	0.00	0.00	0.00	0.00	0.00
2010-11-29	0.00	0.11	0.00	0.00	0.00	0.00	0.51	0.38	0.00	0.00
2011-02-22	0.00	0.00	0.30	0.00	0.00	0.00	0.70	0.00	0.00	0.00
2011-05-16	0.00	0.22	0.00	0.00	0.78	0.00	0.00	0.00	0.00	0.00
2011-08-08	0.00	0.05	0.00	0.00	0.00	0.41	0.00	0.00	0.54	0.00
2011-10-31	0.00	0.01	0.00	0.00	0.45	0.00	0.01	0.53	0.00	0.00
2012-01-23	0.00	0.13	0.00	0.00	0.00	0.00	0.64	0.23	0.00	0.00
2012-04-16	0.00	0.00	0.79	0.00	0.00	0.00	0.00	0.21	0.00	0.00
2012-07-09	0.00	0.22	0.09	0.00	0.69	0.00	0.00	0.00	0.00	0.00
2012-10-01	0.00	0.23	0.00	0.70	0.00	0.00	0.00	0.07	0.00	0.00
2012-12-24	0.00	0.00	0.19	0.00	0.00	0.59	0.22	0.00	0.00	0.00
2013-03-18	0.10	0.00	0.00	0.00	0.65	0.00	0.00	0.24	0.00	0.00
2013-06-10	0.20	0.00	0.00	0.00	0.00	0.00	0.00	0.80	0.00	0.00
2013-09-03	0.16	0.00	0.00	0.00	0.00	0.00	0.71	0.00	0.13	0.00
2013-11-25	0.05	0.00	0.00	0.56	0.00	0.00	0.00	0.39	0.00	0.00
2014-02-18	0.09	0.00	0.00	0.00	0.55	0.00	0.36	0.00	0.00	0.00
2014-05-12	0.05	0.71	0.24	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2014-08-04	0.00	0.23	0.00	0.00	0.00	0.41	0.00	0.00	0.36	0.00

Table 8: Portfolio Weights

	SHY	AGG	HYG	SPY	XLU	XPH	XSD	XRT	XHB	XOP
sum	1.07	2.92	2.92	1.26	4.78	1.65	5.61	5.25	2.54	0.00

Table 9: Column Sum of Weights

Figure 4 shows the cumulative gain for the portfolio, along with the S&P 500 benchmark. This is the  $\hat{G}SPC$  index, not the S&P 500 ETF (SPY). By choosing the index rather than a tradable asset there is a consistent benchmark.




Figure 4: Portfolio Cumulative Return with SP 500 Return

The draw-down chart is shown below in Figure 5






Figure 5: Portfolio and SP 500 Drawdown

The information ratio is:

$$IR = \frac{E[R_p - R_b]}{\sigma} = \frac{E[R_p - R_b]}{\sqrt{\text{var}(R_p - R_b)}} = \frac{\alpha}{\omega} \quad (1)$$

Using the S&P 500 as the benchmark, the information ratio of this portfolio is 0.0016884.

As Figure 4 shows, this portfolio has performance that is very similar to the benchmark.

For taxable funds, the S&P 500 would be a better choice than this portfolio, which is traded every quarter.

A holding period of less than a year-and-a-day will incur short term capital gains taxes. Short-term capital gains are taxed at the income tax rate (e.g., if your tax rate is 28 percent, the short-term capital gains tax will be 28 percent).

In contrast, long-term capital gains (for stocks held at least a year and a day) are 15 percent. The higher short-term capital gains rate incurs a 13 percent penalty compared to simply holding the S&P 500 ETF (SPY).

If a portfolio has performance that is similar to the benchmark (as is the case here), the tax consequences of short term capital gains mean that it would be better to purchase the benchmark and hold it for at least a year.

## **Mean-Variance and Equally Weighted Portfolios**

In this section, for the universe of 10 ETFs (Table 1), an equally equally weighted portfolio is compared to a Markowitz style mean-variance portfolio.

The previous portfolio, shown in Figure 4, is not a classic mean variance portfolio. This portfolio uses the functions shown in Table 7 to estimate the future return. This future return estimate is then used in portfolio optimization.

A Markowitz style mean-variance portfolio uses the mean to estimate the future return. The mean tends to be negatively correlated with the future return, but mean-variance portfolio models do not attempt to correct for this. As Table 5 shows, the mean is not strongly correlated with future returns.

The difference in the estimation of the future return accounts for the difference between the portfolio shown in Figure 4 and the mean-variance portfolio shown in the plot below.




Figure 6: MV Portfolio, Eq. Weighted and SP 500

As Figure 6 shows, the equally weighted portfolio performs better than the mean-variance portfolio and the portfolio in 4. The equally weighted portfolio also has low turnover (resulting in a lower tax rate). If the tax penalty is taken into account, the equally weighted portfolio may be one of the top performing portfolios described in this paper.

## Predicting Returns, Redux

In calculating the forecasted return for the first portfolio, an adjustment factor is used to adjust for negative correlation (in the case of XHB and XPH). The correlation adjustment and the choice of mean function (mean, EMA or OLS) is based on future information (e.g., over-fitting)

## Segmented Linear Model

The more accurately the future return can be estimated, the more likely it is that the portfolio will realize its target return. This section investigates using a segmented linear model to predict future returns. The motivation for this was the idea that a segmented linear model would produce an more accurate estimate of the local mean.

A segmented linear model maps a set of least squares lines through a time series. These fitted lines are the mean through the local distributions.

A segmented linear model is created using the close price. An example is shown below in Figure 7. The return can be calculated from the segmented linear model result.

As with the models discussed previously, one year (52 weeks) of past data is used to develop a prediction of the return, one quarter ahead.




Figure 7: XHB Segmented Linear Model and Close Price

A return sequence is calculated from the fitted values from the linear model. The predicted return is the mean of the fitted values. By using the fitted values from the linear model, there is, in theory, less noise and the price trend is captured.

The table below compares the mean, the segmented linear model mean and the EMA on the segmented linear model values.

	mean	LM mean	EMA
SHY	0.0001	0.0002	0.0005
AGG	0.0009	0.0010	0.0008
HYG	0.0010	0.0008	-0.0007
SPY	0.0028	0.0021	-0.0015
XLU	0.0034	0.0035	0.0015
XPH	0.0070	0.0053	0.0036
XSD	0.0043	0.0029	-0.0130
XRT	0.0010	0.0000	0.0014
XHB	0.0008	0.0001	-0.0018
XOP	-0.0023	-0.0027	-0.0299

The tables below shows the correlation between the values predicted by the mean, EMA and segmented linear model functions with the actual future return.

Mean functions vs. 2-week ahead return:

	SHY	AGG	HYG	SPY	XLU	XPH	XSD	XRT	XHB	XOP
mean	0.1722	-0.0414	-0.0404	0.0421	-0.0128	-0.0378	0.0110	0.0305	-0.0036	-0.0995
EMA	0.0795	0.0327	-0.0496	-0.0046	-0.0658	-0.1176	-0.0235	0.0055	-0.1071	-0.0835
segLM	0.1331	-0.0792	-0.0454	0.0158	0.0079	-0.0191	-0.0254	-0.0268	-0.0334	-0.0710

Mean functions vs. 4-week ahead return:

	SHY	AGG	HYG	SPY	XLU	XPH	XSD	XRT	XHB	XOP
mean	0.2251	-0.0680	-0.0333	0.0876	0.0159	-0.0271	0.0272	0.0267	0.0338	-0.1173
EMA	0.0261	-0.0324	-0.1068	0.0838	0.0147	-0.1650	0.0109	-0.0300	-0.1241	-0.0088
segLM	0.1749	-0.1085	-0.0289	0.0666	0.0277	-0.0127	-0.0232	-0.0154	-0.0293	-0.1077

Mean functions vs. 8-week ahead return:

	SHY	AGG	HYG	SPY	XLU	XPH	XSD	XRT	XHB	XOP
mean	0.3921	-0.0410	-0.0272	0.1413	0.0286	0.0096	0.0743	0.1000	0.1428	-0.0960
EMA	0.2346	0.2141	0.1335	0.1932	0.0297	-0.2804	0.0175	0.1172	-0.0873	0.1568
segLM	0.2886	-0.1062	-0.1031	0.0914	0.0475	0.0039	-0.0239	-0.0542	0.0043	-0.1844

Mean functions vs. 12-week ahead return:

	SHY	AGG	HYG	SPY	XLU	XPH	XSD	XRT	XHB	XOP
mean	0.5057	-0.0223	0.0002	0.1261	0.1518	0.0607	0.1010	0.0690	0.0170	-0.1902
EMA	0.2742	-0.2764	-0.0033	0.3678	0.2586	-0.1758	0.3812	0.3043	-0.1861	0.0960
segLM	0.4105	0.0511	-0.0596	0.0049	0.0493	0.0311	-0.1090	-0.0963	-0.0539	-0.3072

As the comparison point moves out in time (from 2 to 12 weeks), the correlations increase. The EMA is a notably better predictor in most cases than the segmented linear model, the mean or ordinary least squares.

## Looking Back at the Hand Selected Portfolio

The portfolio chosen from the hand selected ETFs by the optimizer for the May to August 2014 time period was:

1	2
HYG	0.24
XLU	0.76

The optimization was not constrained to limit portfolio weight and there were only nine ETFs in the ETF universe.

This portfolio may suffer from over-fitting. In selecting the return forecast functions and the correlation adjustment, future knowledge is used to select the function with the best historical correlation and to select the adjustment factor (to adjust for negative correlation). Small changes in the mean functions and the correlation adjustment result in large changes on the portfolio result. This suggests that these factors are not stable.

Another problem is that the correlation adjustments for the various functions are estimated from the distributions of only a few assets.

## Portfolio Redux

The rest of the paper investigates portfolios that have better performance than the portfolio shown in Figure 4, which was constructed from the hand selected set of ETFs.

## Asset Selection

Portfolio returns are extremely sensitive to the choice of assets. Quantitative techniques cannot turn lead into gold. A portfolio algorithm can only produce returns within its universe of assets. For example, the maximum

gain-loss and maximum Sharpe ratio algorithm produced extremely poor returns with an initial set of ETFs. Much better returns were produced with the etf.com list, discussed below.

The literature on asset selection for stocks goes back at least as far as the first edition of Graham and Dodd's 1934 book *Securities Analysis*. A variety of factors have been proposed for selecting the stocks for a portfolio.

ETFs are collections of stocks and other assets (like futures contracts) that can be traded without tax consequences. This makes the analysis of ETFs more complicated than single stocks. ETFs also have a wide variety of assets and strategies. Their complex nature makes ETFs more difficult to select on an empirical basis.

The portfolios select ETFs on the basis of past performance. The problem with always looking at the past, without any tools to forecast future performance, is that sometimes the desired performance is also in the past.

## Selecting ETFs

In the case of stocks, the universe can be selected from the stocks that make up an index, like the S&P 500. With ETFs there is no canonical choice for the ETF universe.

Originally a list of ETFs was constructed from a list from the Morningstar web site. The Morningstar ETF filter is difficult to use and does not allow data to be exported to a CSV file. Morningstar has also placed a number of features behind a pay wall.

I later discovered that a better source of ETF information is the etf.com web site. This site allows ETFs to be selected on the basis of minimum capitalization (among other criteria). The data can be easily exported in CSV format.

The "settings" used in the etf.com screen are:

- Efficiency: 80 and above
- AUM: >\$10 million
- 3 Mo Return: >0 percent

- 1 Year Return: >0 percent
- Fund Closure Risk: Low Risk

In November of 2014 this resulted in a universe of 436 ETFs.

Most ETFs have been constructed in the last few years (relative to 2014). To allow sufficient data for back testing, a start date of 2007-04-04 was used. Out of the 436 ETS, there were 217 that had an inception date in this time frame.

## Asset Selection and Risk

If there is sufficient data, assets can be selected from an asset universe using mean-variance optimization. Mean-variance analysis relies on the calculation of a covariance matrix. In most cases, we do not have sufficient past data, relative to the asset universe. This creates a significant problem when it comes to estimating the covariance matrix. With just a year of past data (52 samples), the ETF universe is larger than the sample data. The resulting covariance matrix may not be full rank and may have a great deal of error (see *Estimating High Dimensional Covariance Matrices and its Applications* by Bai and Shi (2011)).

In *Honey, I Shrunk the Sample Covariance Matrix* (2003) Ledoit and Wolf describe how "shrinkage" can be used to construct a more accurate covariance matrix. Ledoit and Wolf's shrinkage algorithm is implemented by the R function `cov_shrink` in the `tawny` package. This package is used to estimate the covariance matrix for this portfolio.

Even if Ledoit and Wolf's shrinkage method is used, such a large mismatch between the number of assets and the number of historical samples may still yield an inaccurate covariance matrix. By pre-filtering assets to yield an asset universe that is closer to the number of historical data points, a more accurate covariance matrix may be achieved.



## Asset Selection using the Gain-Loss Ratio

In the paper *Gain, Loss and Asset Pricing* by Bernardo and Ledoit (1996) the authors propose an asset performance metric that improves on the Sharpe Ratio by only penalizing for the down side loss. This metric also has the advantage that, unlike the Sharpe Ratio, the sample standard deviation does not have to be estimated (standard deviation has a lot of error for small samples). Unlike their later article on covariance matrix shrinkage, this article is not exactly a paragon of clarity. However, the gain-loss equation is nicely summarized in the book *Practical Risk-Adjusted Performance Measurement* by Bacon, page 110 (available on Google Books). The equation is:

$$\text{gain-loss ratio} = \frac{\frac{1}{n} \times \sum_{i=1}^n \max(r_i, 0)}{\frac{1}{n} \times \sum_{i=1}^n \max(0 - r_i, 0)}$$

Note that bond ETFs tend to have high gain-loss ratios, since they tend to have a small number of losses and lots of (small) gains. Although bond ETFs will have a high gain-loss ratio, their return may be low.

The plot below shows the distribution of gain-loss ratio values for the time period 2007-04-09 to 2008-03-31. The back-test time period includes the crash of 2007-2008, so there are probably an unusually large number of gain-loss ratio values that are less than one (e.g., more losses than gains).

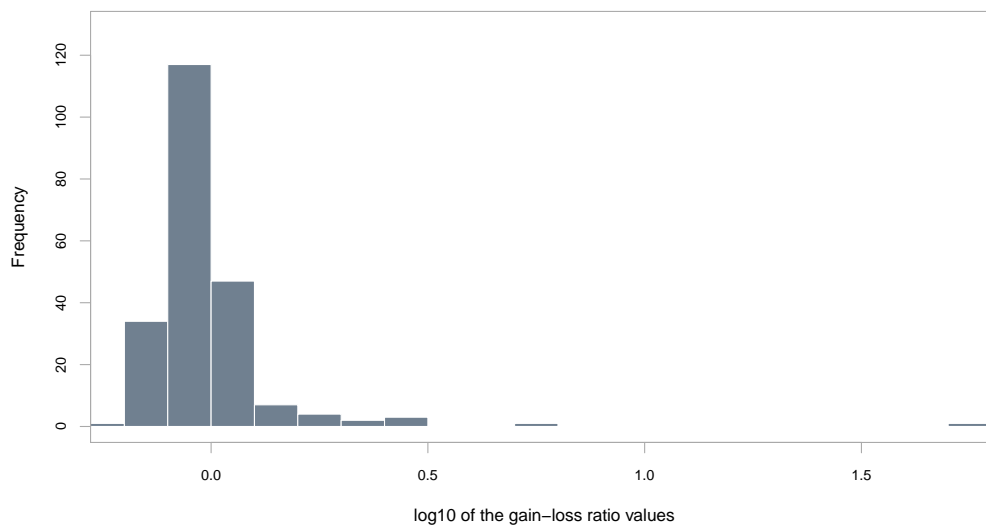


Figure 8: Histogram of the log10of the gain-loss ratio

## Back-testing Portfolio Algorithms

Back-testing uses a rolling 1-year window that moves forward in 12-week steps.

For back-testing the asset universe is limited to ETFs that span the entire back test period. This limitation exists for the sake of simplicity (a new asset universe doesn't have to be calculated every quarter).

## Maximum Return Portfolios

The algorithm for calculating a portfolio that targets the maximum forecasted return is:

1. From 52 weekly returns, calculate the Sharpe ratio or the gain-loss ratio for all assets.
2. Sort the assets in decreasing order on the basis of the Sharpe or gain-loss ratio result.
3. For each ETF, remove the other ETFs that are highly correlated ( $\rho \geq 0.9$ ) This helps avoid a portfolio that is composed the same type of ETF (e.g., energy, technology, etc).
4. Select the top N assets from the set that remains after correlation filtering.
5. Calculate the forecasted return for the selected assets. Return is forecasted using either the EMA function.
6. Use shrinkage to estimate the covariance matrix based on 52 weekly returns
7. Pick the target return from forecasted future returns:  
 $r_{target} = max(\mathbf{r}_p) - (max(\mathbf{r}_p) \times frac)$ , where  $\mathbf{r}_p$ . This target return is near the maximum projected return for the portfolio and  $frac$  is usually between 0.1 and 0.3. For this sets of ETFs,  $frac = 0.3$
8. Using numeric optimization, calculate the optimal portfolio for the selected ETFs

In the Sharpe ratio the risk free rate is subtracted from the return. Here a modified Sharpe ratio is used:

$$\text{Sharpe ratio}_{modified} = \frac{\mathbb{E}[\mathbf{r}]}{\sigma_r}$$

The back-test result, for the maximum gain-loss and Sharpe ratio filtered portfolios, is shown below in Figure 9.




Figure 9: Plot of Portfolio Redux Performance

The Maximum gain-loss and Sharpe ratio portfolios are extremely sensitive to the choice of assets. With an earlier set of ETFs this portfolio lost about 60 percent of its value in the crash of 2008-2009. By the end of the back-test period in 2014 this portfolio had only recovered about 20 percent of its value, while the S&P 500 had recovered its loss, plus a 40 percent gain.

In contrast, with the current set of ETFs, both of maximum gain-loss and Sharpe ratio portfolios significantly out perform the S&P 500 benchmark.

## Optimized Portfolios and Diversification

The initial portfolio, before optimization, consists of 40 ETFs. The plot in Figure 10 shows the number of ETFs in the optimized portfolio. In some cases the portfolio consists of only two ETFs.




Figure 10: Number of ETFs per Portfolio

One of the criticisms of optimized portfolios is that they concentrate the portfolio in a few assets, violating the portfolio construction dogma of diversification. While reciting point, the critic will point out that diversification reduces risk, which can be shown mathematically.

The maximum Sharpe and gain-loss ratio portfolios are given a target return (the maximum forecasted return, using the EMA function). The optimizer will find the combination of assets from the input portfolio that produces that return with the lowest risk. The optimizer can be given constraints that limit the weight for any asset (e.g., no asset can make up more than 10 percent of the portfolio). Constraining the portfolio in this way usually reduces the possible portfolio return. To arrive at a portfolio where the optimizer can find a solution, a lower target return will be necessary. The portfolio that results will have a lower return than the unconstrained portfolio, but also a lower risk.

One argument for a diversified portfolio is that it performs better (and has lower losses) in changing market conditions. The longer the portfolio is held,

the more important this is.

The design of the portfolios in this paper take a different approach. Changing market conditions are adapted to by re-balancing the portfolios with a new set of assets.

## Using `portfolio.optim()`

The portfolio optimization above uses `solve.QP`, whose arguments can be confusing. In this section `portfolio.optim` is applied to check that the results are the `solve.QP` results.

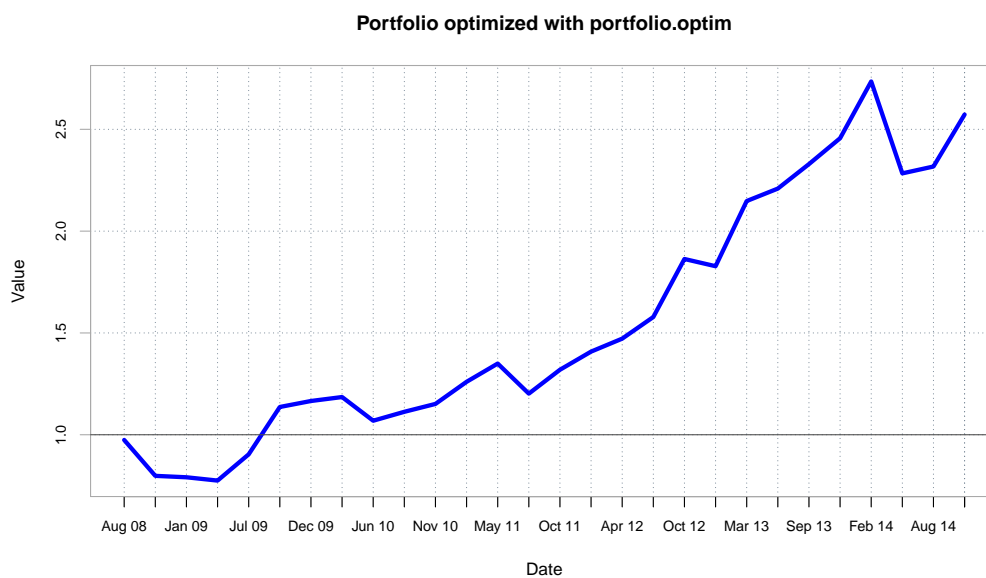


Figure 11: `portfolio.optim()` optimized Portfolio (gain-loss filtered)

This is the same result as the gain-loss and Sharpe ratio portfolios optimized using `solve.QP`. The application of `solve.QP` appears to be correct.

## Equal Weighted Portfolios

Comparing an equally weighted portfolio to an optimized portfolio provides a useful perspective on whether the optimization resulted in a portfolio with better performance.

An equally weighted portfolio is examined in this section.

1. From 52 weekly returns, calculate the gain-loss or Sharpe ratio for all assets.
2. Sort the assets in decreasing order on the basis of the gain-loss ratio
3. For each ETF, remove the other ETFs that are highly correlated ( $\rho \geq 0.9$ )
4. Pick the top  $n$  assets  $n = \min(\text{numAssets}, 20)$
5. Set the portfolio weights to  $\frac{1}{n}$




Figure 12: Plot of Equal Weighted Portfolio (gain-loss filtered)

The equally weighted portfolio out performed the benchmark portfolio, but with a lower loss during the market crash (20 vs. a 60 percent loss).

## Efficient Frontier Portfolios

The maximum return portfolio targets a return close to the maximum forecasted return.

The portfolios constructed in this section start in the same way: a set of assets are selected using the Sharpe ratio or the gain-loss ratio. However, rather building a single portfolio for a target return, an efficient frontier is calculated for a range of returns. The portfolio on the efficient frontier with the highest Sharpe or gain-loss ratio is selected.

The previous portfolios used the EMA function to estimate the future return. The portfolios below use both EMA and the mean to estimate the future return.

## Efficient Frontier Portfolios

This section explores the construction of efficient frontier portfolios. The result portfolio is selected on the basis of the highest gain-loss or Sharpe ratio for the portfolio.

For each back-test step:

1. From 52 weekly returns, calculate the gain-loss ratio for all assets.
2. Sort the assets in decreasing order on the basis of the gain-loss ratio
3. For each ETF, remove the other ETFs that are highly correlated ( $\rho \geq 0.9$ )
4. Pick the top  $n$  assets  $n = \min(\text{numAssets}, 40)$
5. Estimate the future returns,  $\mathbf{r}_p$
6. Find the efficient frontier for the long-only portfolios using target returns in the range  $\min(r_p) \cdots \max(\mathbf{r}_p) - (\max(\mathbf{r}_p) \times 0.10)$



7. For each portfolio, return the estimated future return and the portfolio standard deviation.
8. Find the maximum Sharpe Ratio portfolio.

The maximum Sharpe Ratio portfolio chosen from the efficient frontier has no loss during the 2008 crash and shows steady appreciation. However, the terminal value of the portfolio is less than the S&P 500 benchmark portfolio.




Figure 13: Plot of the maximum Sharpe Ratio Portfolio

The IR of the portfolio is  $-0.0650078$ . The draw-down plot is shown below.




Figure 14: Max Sharpe Portfolio and SP 500 Drawdown

For the risk averse, this portfolio seems promising for this set of assets.

The Sharpe Ratio has the disadvantage of penalizing for upside volatility, in addition to downside volatility. Instead of choosing the efficient frontier portfolio with the maximum Sharpe Ratio, this algorithm chooses the portfolio with the maximum gain-loss ratio.




Figure 15: Plot of the maximum Gain-Loss Ratio Portfolio

The IR of this portfolio is  $-0.1207773$ .

For this set of assets, the efficient frontier portfolio selected by its gain-loss ratio performs worse than the portfolio selected using the Sharpe ratio.

Here again it is worth noting that the portfolio characteristics depend heavily on the assets. For another set of assets the Sharpe and gain-loss portfolios had similar performance.

## Return Ranked Portfolios

In this section portfolios are constructed by ranking the assets by their mean return. The implication of this approach is that the mean return is assumed to be a predictor of future return (an idea that was resisted in previous portfolios).

The previous section shows results for portfolios constructed from assets that

are ranked either by their Sharpe ratio or their gain-loss ratio. These ratios can be independent of the magnitude of return. For example, an assets with a small return but an even smaller standard deviation will have a high Sharpe ratio. In the case of the gain-loss ratio, the actual magnitude of the gain or loss doesn't matter.

Portfolios constructed using the Sharpe or gain-loss filters can be heavily weighted toward bonds. While these portfolios have lower risk, they may be overly conservative. This section examines portfolios where the assets are ranked on the basis of their estimated future return.

## Return Ranked, Maximum Sharpe Ratio Portfolio

- Rank the assets by their mean return
- Filter out assets that have a high paired correlation with a reference asset.
- Select  $N$  assets for the portfolio.
- Use the mean return as an estimate of the future return <sup>2</sup>
- Calculate the efficient frontier for a range of target returns.
- Select the portfolio with the highest Sharpe ratio.

---

<sup>2</sup>Previous portfolios use EMA to estimate the future return. In theory this is a better estimate than the mean because it has a higher correlation with future return. However, this doesn't seem to be true here. Using the mean to estimate future return is a result of the observation that this produces a portfolio with better performance. However, in doing this there is a real danger of over-fitting.




Figure 16: Plot Return Filtered Max Sharpe Ratio Portfolio

By ranking the portfolio assets by their mean return, higher return assets are selected, which may completely omit assets like bond funds with lower return and risk.

The next portfolio mixes 90 percent assets that are selected by return ranking and 10 percent selected by gain-loss ranking (these percentages were selected by observing back-test performance).




Figure 17: Plot of the Maximum Return + 10 percent Gain-Loss Ratio Portfolio

The best performing portfolios are the portfolios that target the maximum forecasted return from assets that were filtered using the gain-loss or Sharpe ratio (see Figure 9).

Unfortunately these portfolios are not without their problems. In constructing the optimized portfolio, a target return near the maximum asset value is used. This can lead to selecting assets that perform very badly in market crashes. Although these portfolios perform well with the current set of assets, there was very poor performance with a previous set of ETFs.

Of the ranked portfolios that do not have significant draw-down in the market crash, best performing portfolio seems to be the maximum Sharpe ratio portfolio (Figure 16) using the mean to calculate the future return and filter the portfolio set. This performance seems to be relatively independent of the set of assets.

# Actual Performance

## August 28, 2014 - November 28, 2014

This section looks back the performance of the portfolio selected for the period August 28 to November 28, 2014. The portfolio is shown in the table below. The portfolio is a maximum Sharpe ratio portfolio (see Figure 16).

GSY	DIV	PZA	MINT	IYW	CSM	BAB	NIB	FBT	GVT
0.3410	0.0130	0.0840	0.4720	0.0240	0.0040	0.0340	0.0110	0.0100	0.0070

This plot shows weekly continuously compounded returns. This return is unrealistically high, since the portfolio is not continuously compounded. This plot does give an idea of what the relative performance of the portfolio vs. the benchmark would be.




Figure 18: Aug. 28, 2014 portfolio vs. SP 500

## Current Investment Portfolio: December 1, 2014

The current portfolios are constructed for the current 1-year period. This provides a larger universe of ETFs than the back-test. Of the 439 ETFs there were 431 which had a year of past data.

Two portfolio models were considered for the current investment period (e.g., December 1, 2014):

- Maximum return target portfolio. From August 2008 to November 2014 the maximum return portfolios in 9 has a return of almost 150 percent. This portfolio did have a loss during the 2008-2009 market crash. This portfolio is also extremely sensitive to its ETF universe composition.
- Efficient frontier maximum Sharpe ratio portfolio (Figure 16). This is a more conservative portfolio. This portfolio had no loss during the market crash, although in the same period it had a very small return. During the investment period it had a return of almost 40 percent. This portfolio model is less sensitive to its ETF universe and had similar performance with past ETF universes.

### Maximum return target portfolio

sym	weight	name
FBT	0.4400	First Trust NYSE Arca Biotechnology
PJP	0.0900	PowerShares Dynamic Pharmaceuticals
RYU	0.4300	Guggenheim S&P Equal Weight Utilities
INY	0.0500	SPDR Nuveen Barclays New York Municipal Bond

The correlations between the assets are:



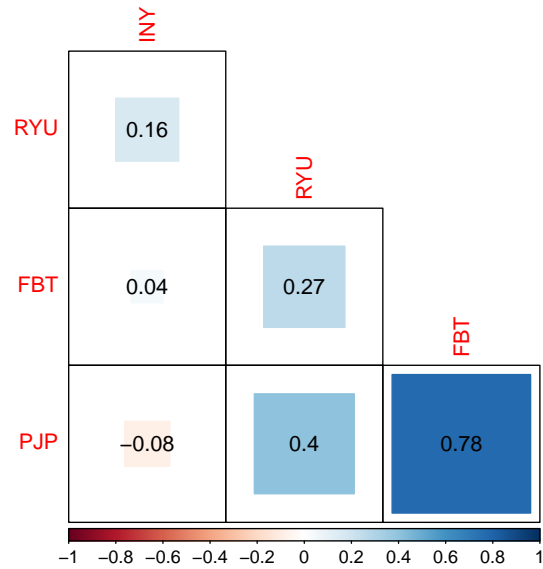


Figure 19: Correlation of Max Sharpe Ratio Current Portfolio

## Efficient Frontier Maximum Sharpe Ratio Portfolio

sym	weight	name
FBT	0.0500	First Trust NYSE Arca Biotechnology
INDA	0.0310	iShares MSCI India
XPH	0.0420	SPDR S&P Pharmaceuticals
IHF	0.1310	iShares U.S. Healthcare Providers
PEY	0.0120	PowerShares High Yield Equity Dividend Achievers
DIV	0.2020	Global X SuperDividend U.S.
BABS	0.5080	SPDR Nuveen Barclays Build America Bond
COW	0.0230	iPath Dow Jones-UBS Livestock Total Return ETN

The correlation between the portfolio assets is shown below.

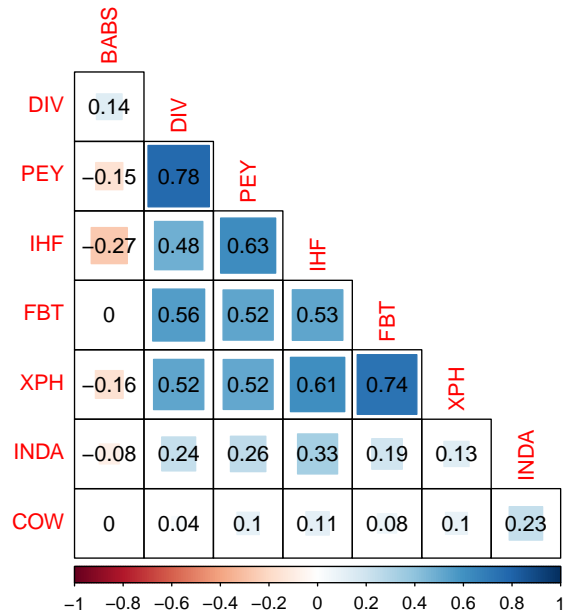


Figure 20: Correlation of Max Sharpe Ratio Current Portfolio

For smaller portfolios, asset weights of only a few percent can result in trading costs that consume a significant fraction of the asset return (at least for retail brokers like Charles Schwab). Portfolio optimization can allocate a portfolio that has no more than a certain percentage for each asset. But there is no easy way to allocate at least X percent per asset.

To address this issue, the weights for assets below a particular threshold can be proportionally reallocated to the other assets. This will change the risk return profile for the portfolio, but it should not be a huge change.

The portfolio, rounded to reallocate any asset with a weight of 5 percent or less, is shown below.

IHF	DIV	BABS
0.1600	0.2400	0.6000

## And the winner is...

Although the maximum target return portfolio has more risk, it's historical returns are hard to resist. We'll see in early March 2015 how this worked out. The portfolio is shown below:

sym	weight	name
FBT	0.4400	First Trust NYSE Arca Biotechnology
PJP	0.0900	PowerShares Dynamic Pharmaceuticals
RYU	0.4300	Guggenheim S&P Equal Weight Utilities
INY	0.0500	SPDR Nuveen Barclays New York Municipal Bond

## Reproducibility

This is reproducible research. This paper is compiled from a combination of executable R code and  $\text{\LaTeX}$ text. All of the tables and plots in the paper were generated by executing the R code that makes up the source of the paper. This paper is published along with its source and the ETF lists that were used. All of the ETF market data used in the paper was downloaded from `finance.yahoo.com`. The reader can examine the R code in the paper source to see how each of the plots was generated. The paper can be rerun using Knitr and R to regenerate the paper.