Momentum Based Balancing for the Diversified Portfolio – Abstract

February 28, 2014

National Association of Active Investment Managers (NAAIM)

Wagner Award 2014

Mark Virag

Virag Capital, Inc.

mark.virag@viragcapital.com

Background

The traditional purpose of portfolio balancing is to maintain a predetermined mix of investments, offsetting the differing rates of growth among instruments in the portfolio by periodically redistributing capital to the slower growing instruments. Moving capital from strong performers to weak ones may seem counterintuitive, but this approach is still considered best practice by many portfolio managers.

Depending on the portfolio selection, the time horizon, and the allocation mix, the traditional approach may produce superior risk adjusted returns. There are many studies that support this approach, particularly with regard to reducing portfolio volatility.

The present paper follows a contrary thread in the literature, which holds that the manager should actively change the allocation mix in response to changing market conditions. Various rules for active balancing have been proposed. One challenge with this approach, in contrast to the traditional one, is arriving at rules which are comparably clear and straightforward. The rules must be tested empirically, especially with attention to volatility, and also supported by theory. Finally, we would like to unseat the primacy of portfolio selection, and subordinate it to the balancing rules.

Methods

In this paper, we present the use of monthly price movements as the basis for active balancing. We test the first model portfolio using three alternative rule sets on this basis, and we show that all three produce superior risk adjusted returns relative to an equal balance model, and relative to the general market.

For the model portfolio, we use an assortment of Exchange Traded Funds (ETFs) chosen to exhibit the desired broad diversification. For the benchmark, we rebalance monthly to have equal sums invested in each of the instruments, which are 30% bonds, 40% stocks, and 30% to cash, commodities, and real estate. We also benchmark against **SPY** as a proxy for the general market and **TIP** as the risk free rate of return. The paper supplies a formal definition of portfolio diversity with reference to average inter-item correlation, and also with reference to practical and fundamental characteristics.

For the alternative balancing rules, we present three models which allocate capital pro-rata among the instruments based on each one's performance in the prior month. The three are more or less aggressive depending on their objectives for portfolio stability, i.e., the concentration of capital in many or few of the instruments in the model portfolio. We describe each model in sufficient detail to be replicated. The balancing rules are straightforward.

Results

We compare the performance of each model using total return, compound annual growth, annualized standard deviation, and drawdown. We use the Sharpe ratio for risk adjusted return, and we also consider the performance of the individual instruments. We find that the pro-rata models handily outperform the benchmarks on most metrics.

This result is supported by the theory that, while monthly price action is not predictive for an individual instrument, relative rankings among a portfolio of widely diversified instruments is predictive. We use a regression study and a hypothesis test to prove the theory for a specific portfolio, and we discuss the application of inter-item correlation to portfolio selection. This addresses our goal of using quantitative measures to choose instruments that suit the balancing rules.

Each model looks back only to the prior month because it is allocating for one month, and these periods should match. Our goal is to demonstrate that price action alone is predictive over the short term, without recourse to a longer history or historical indicators. The same could be said of balancing on a quarterly period. We test the model's sensitivity to this assumption, and find that using a three month lookback period enhances the model's risk adjusted performance. We show that this is because the longer lookback increases the determination of trailing returns based on relative rank.

We also test for sensitivity to correlated returns, and find that the theory and the performance do not hold up where the instruments are insufficiently diverse. We demonstrate this by repeating the experiment using a second model portfolio having high inter-item correlation. This portfolio consists of the nine sector funds commonly used in rotation models. In this case, the dynamic model outperforms the benchmark by only a narrow margin. Relative rank proves not to be predictive, owing to close correlation among monthly price movements of the sectors.

Conclusion

We conclude that the relative ranking theory is a sound basis for portfolio balancing, even with a short lookback period, provided that the instruments in the portfolio are sufficiently diverse. An average inter-item correlation of 35% is low enough, and 55% is too high. The latter figure represents the nine SPDR sector funds, which will be familiar to most readers. For monthly balancing, a lookback period of one month is sufficient, and three months is better.

Momentum Based Balancing for the Diversified Portfolio

February 28, 2014

National Association of Active Investment Managers (NAAIM)

Wagner Award 2014

Mark Virag

Virag Capital, Inc.

mark.virag@viragcapital.com

Overview

This paper presents a novel approach to portfolio balancing, based on short term price movements among instruments in a broadly diversified portfolio. We test a model portfolio using three alternative rule sets on this basis, and we show that all three produce superior risk adjusted returns relative to an equal balance model, and relative to the general market. We construct the portfolio to minimize inter-item correlation, so that its risk and return characteristics are determined primarily by the balancing rules.

For the benchmark, we rebalance monthly to have equal sums invested in each of the instruments. We also benchmark against **SPY** as a proxy for the general market and **TIP** as the risk free rate of return.

We compare the performance of each model using total return, compound annual growth, annualized standard deviation, and drawdown. We use the Sharpe ratio for risk adjusted return, and we consider the performance of the individual instruments. We also consider practical aspects of the balancing approach:

- Limited number of rebalancing events
- Stability of holdings over time
- Simple rules, suitable for a nonprofessional

- Rules for going to cash

The three variations outperform the benchmark, and we present a statistical proof of why "momentum" is effective in this context. We use a regression study and a hypothesis test to prove the theory for a specific portfolio, and we discuss the application of inter-item correlation to portfolio selection.

We then test the model's sensitivity to correlation, and its sensitivity to the lookback period. Finally, we discuss the data series used in the study, and implementation details.

Four Models

We begin with a traditional[1] balanced portfolio consisting of 30% bonds, 40% stocks, and 10% each in cash, commodities, and real estate. We enforce an equal allocation by rebalancing once a month. Thanks to the proliferation of ETFs, an investor can simply buy equal amounts of the following (leaving 10% in cash).

SPY - SPDR S&P 500 ETF Trust

EFA - iShares MSCI EAFE Index Fund

IEF - iShares Barclays 7-10 Year Treasury Bond Fund

TIP - iShares Barclays TIPS Bond Fund

EEM - iShares MSCI Emerging Markets Index

IWM - iShares Russell 2000 Index

XLB - Materials Select Sector SPDR

IYR - iShares Dow Jones US Real Estate

TLT - iShares Barclays 20+ Yr Treasury Bond

Selection criteria for the portfolio are discussed in a later section. The chart below shows how the nine funds have performed individually over the last ten years.

---

[1] Vanguard Group, "Best practices for portfolio rebalancing," July 2010.

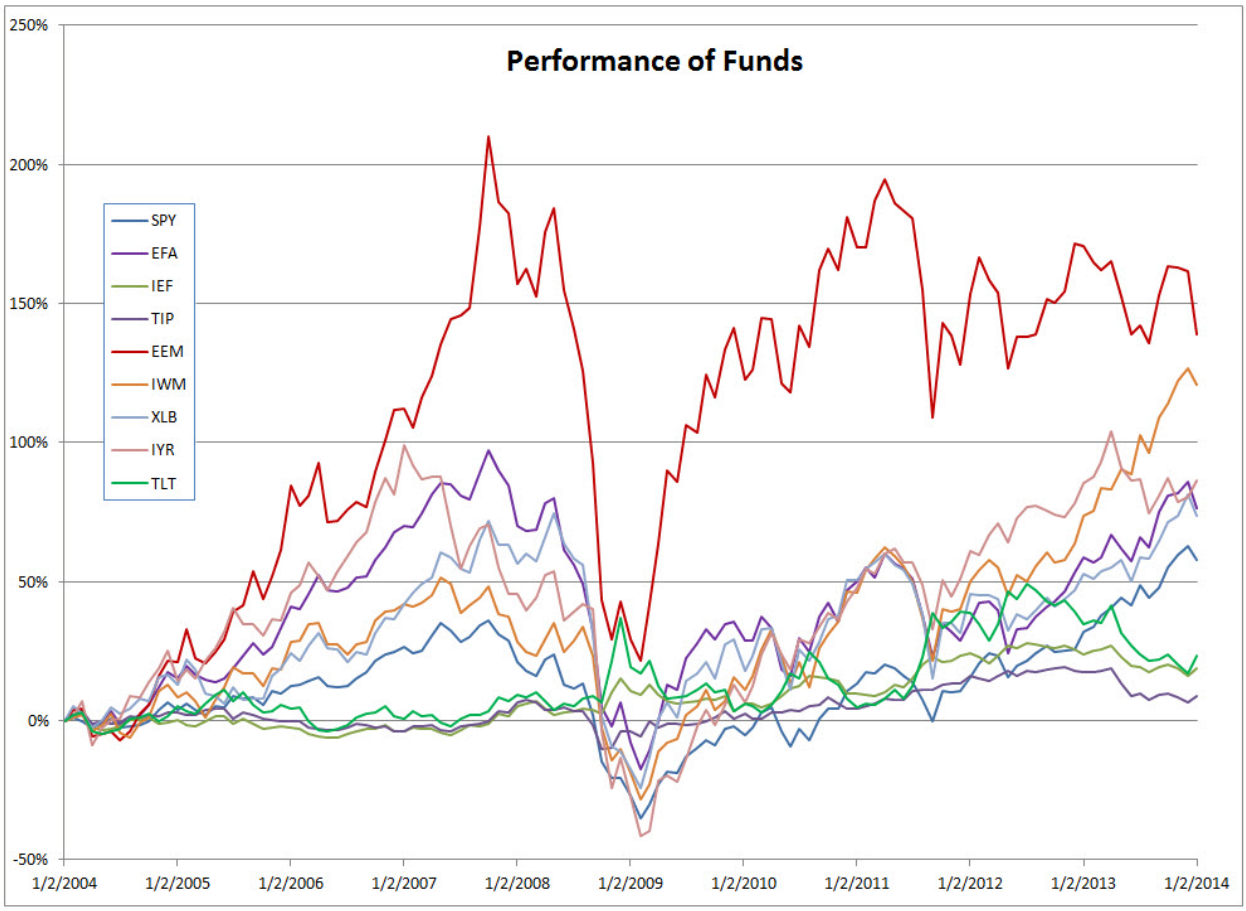

Taking **TIP** as the risk-free rate of return, here are the performance metrics:

	SPY	EFA	IEF	TIP	EEM	IWM	XLB	IYR	TLT
Total Return	57.9%	76.3%	19.1%	8.8%	139.1%	120.6%	73.8%	86.6%	23.6%
Annualized Growth	4.7%	5.8%	1.8%	0.8%	9.1%	8.2%	5.7%	6.4%	2.1%
Standard Deviation	14.6%	19.0%	6.6%	6.7%	24.9%	19.6%	21.0%	25.2%	13.6%
Sharpe Ratio	0.26	0.26	0.14	0.00	0.33	0.38	0.23	0.22	0.09

Emerging markets were strong but volatile over the period. Considering volatility, the Russell was a better investment. Below is the

balanced portfolio versus **SPY**. We start with $1,000 in each of the nine ETFs, plus $1,000 in cash. Each month, we rebalance the total, one tenth into each category. This is a natural, intuitive balancing rule. We'll call it model #1, or "equal balance."

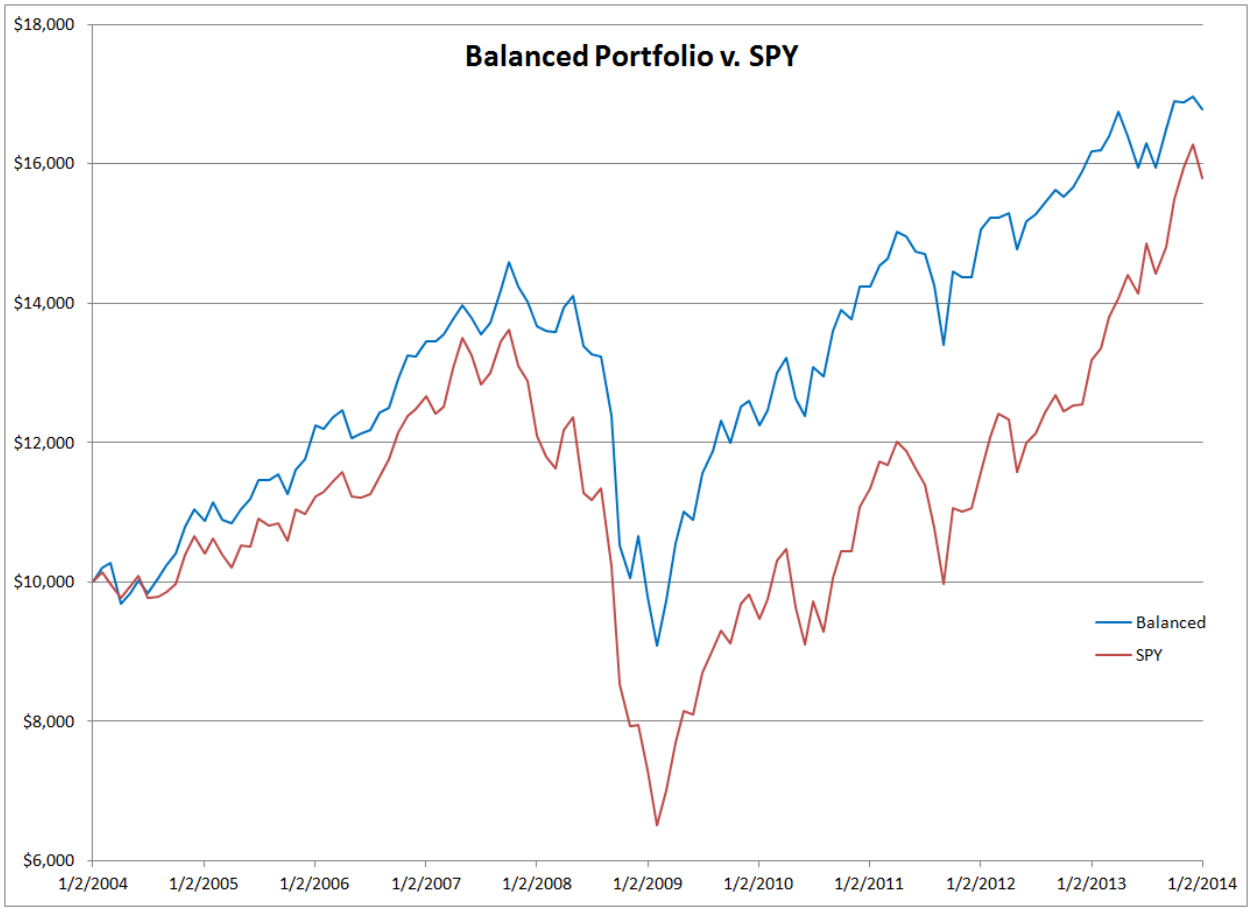

Model #1 has an annualized growth rate of 5.3% and standard deviation of 11.2%. It is less volatile than all but the bond funds, and its 0.40 Sharpe ratio dominates all of the individual ETFs.

As Loeb observes, equal balance has the effect of moving capital from issues that are performing well to those that are not[2]. In the next model, we rebalance pro-rata according to how well each fund has done in the prior month.

In the first month we start with equal allocations to the nine ETFs, and we make a 2.3% return. Model #2 does not hold cash.

SPY	EFA	IEF	TIP	EEM	IWM	XLB	IYR	TLT	Avg.	Total
1.4%	2.3%	1.3%	2.1%	3.9%	1.0%	5.2%	1.8%	1.9%	2.3%	20.9%

For the next month, we allocate according to how each fund performed as a percentage of the total. This hypothetical "total" figure is only used for the pro-rata calculation. Here is the result:

SPY	EFA	IEF	TIP	EEM	IWM	XLB	IYR	TLT
6.5%	11.0%	6.3%	10.3%	18.9%	4.6%	24.9%	8.6%	8.9%

In this month, **XLB** did best, so it is allocated the most capital going into the next month.

XLB Allocation = 24.9% = 5.2 / 20.9

[2] Loeb, Gerald, "The Battle for Investment Survival" (Wiley Investment Classics, 2007)

Model #2 is the "winners only" model. We repeat this procedure each month, allocating capital only to ETFs with positive returns, and in proportion to the size of those returns relative to the others. Since the return to cash is always zero, this model never allocates to cash. Here is how it performs:

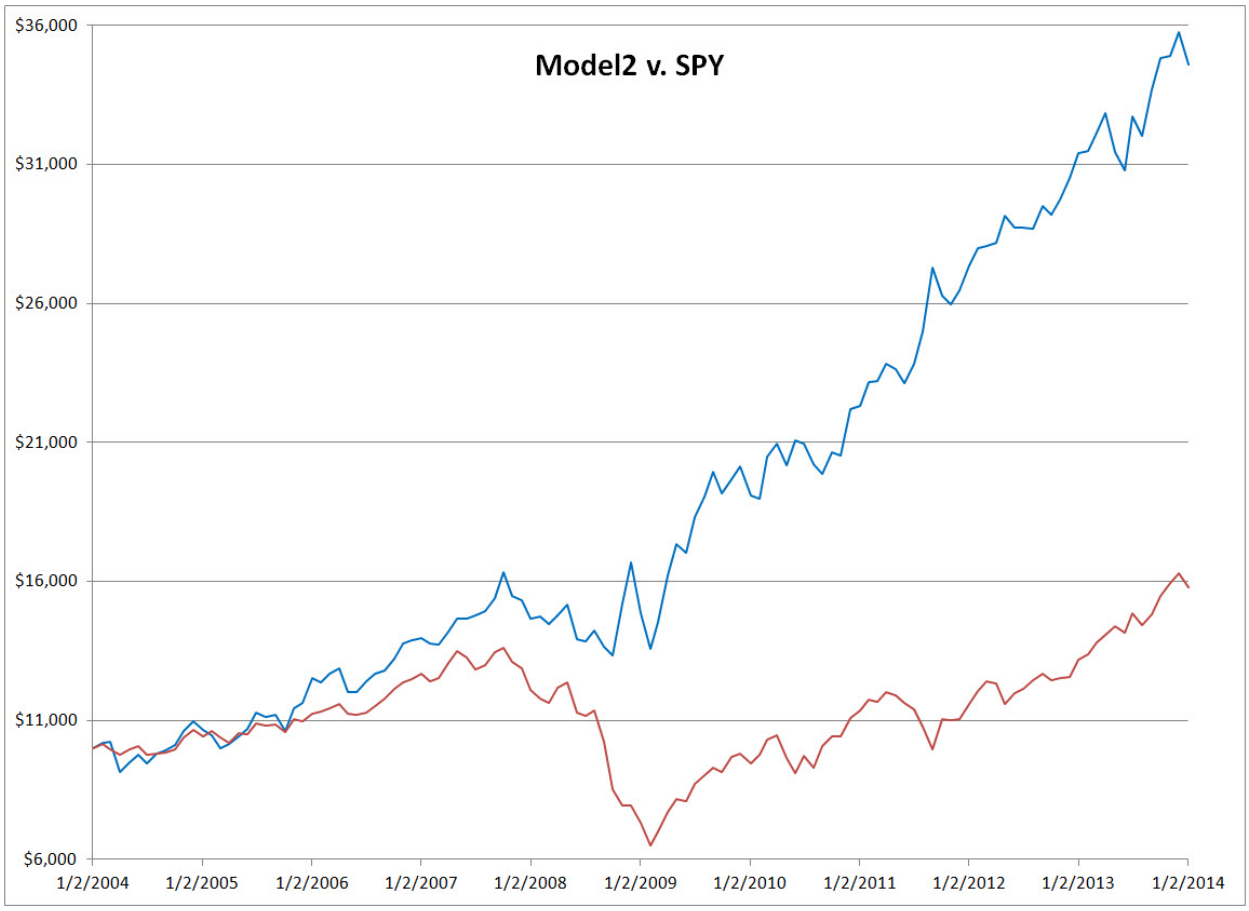

Model #2 has an annualized growth of 13.2%, dominating all the individual funds, and with less volatility than SPY. Its Sharpe ratio is 0.86.

Results for this model are impressive, but it has some drawbacks. In some months all issues lose money, leaving no choice but to reuse the prior

month's rankings. The model also cannot go to cash. Finally, it produces an erratic mix of issues from month to month.

The chart below shows the portfolio mix over the course of a representative year, 2005.

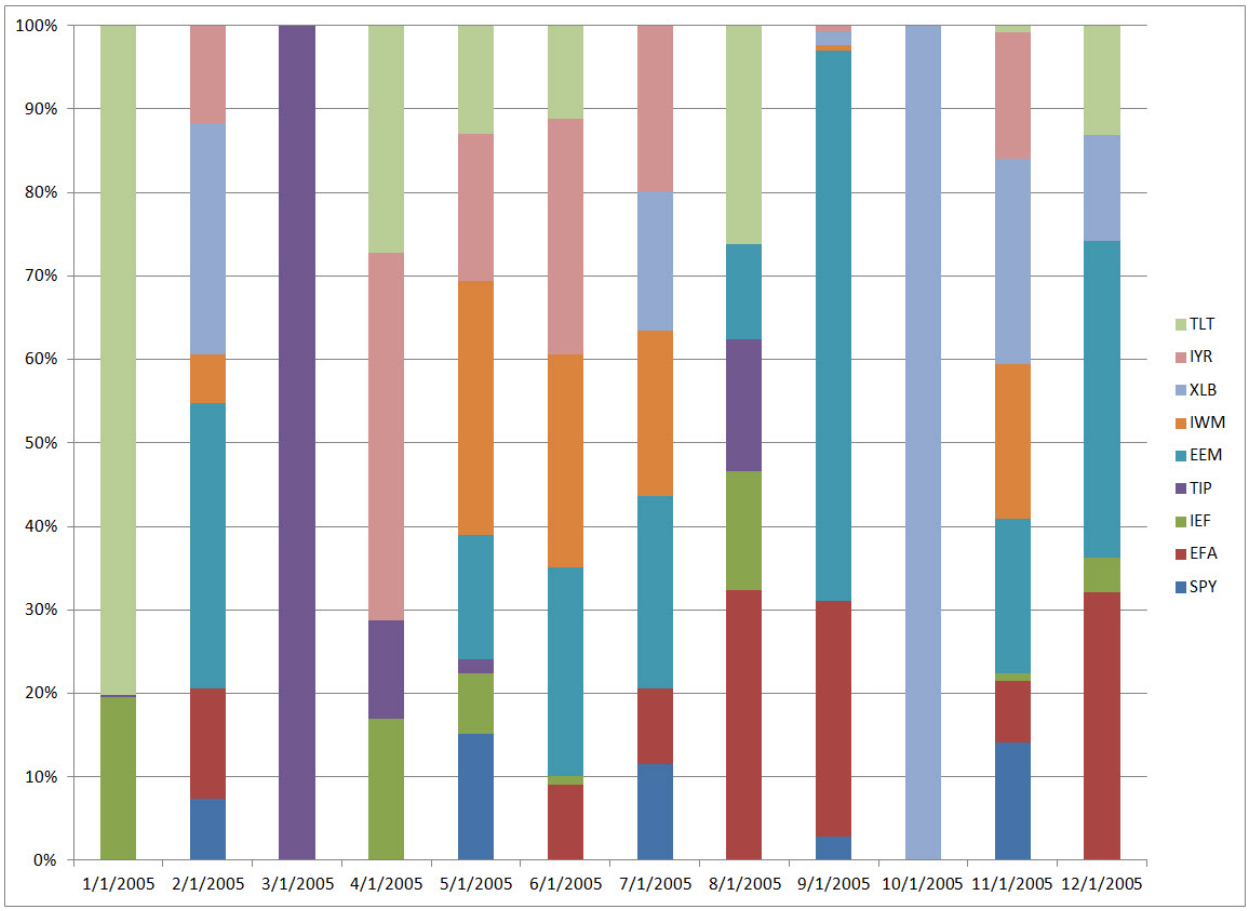

This is too much churn for most investors. For the next model, we keep the pro-rata concept, but we resolve to stay in nine of the ten issues (including cash) every month. We do this by baselining all issues to the month's worst performer.

For example, in a month (March 2004) where the returns are:

SPY	EFA	IEF	TIP	EEM	IWM	XLB	IYR	TLT	Cash
-1.7%	0.1%	1.0%	1.7%	0.6%	1.0%	-2.6%	5.4%	1.0%	0.0%

We find the biggest loss, 2.6%, and add that amount to each return:

SPY	EFA	IEF	TIP	EEM	IWM	XLB	IYR	TLT	Cash
1.0%	2.7%	3.6%	4.3%	3.2%	3.6%	0.0%	8.0%	3.6%	2.6%

Then, we repeat the pro-rata allocation as before. Note that **cash**, which always returns zero, receives an allocation when at least one ETF makes a loss. In this example, **XLB** is the biggest loser, and so the portfolio will hold no shares of **XLB** in the next month. If all issues make gains, then **cash** is weakest, and the model goes fully into ETFs.

This is the "drop one loser" model. It's not as erratic as model #2, but it also doesn't perform as well. Its annualized growth of 8.1% is less than **IWM**.

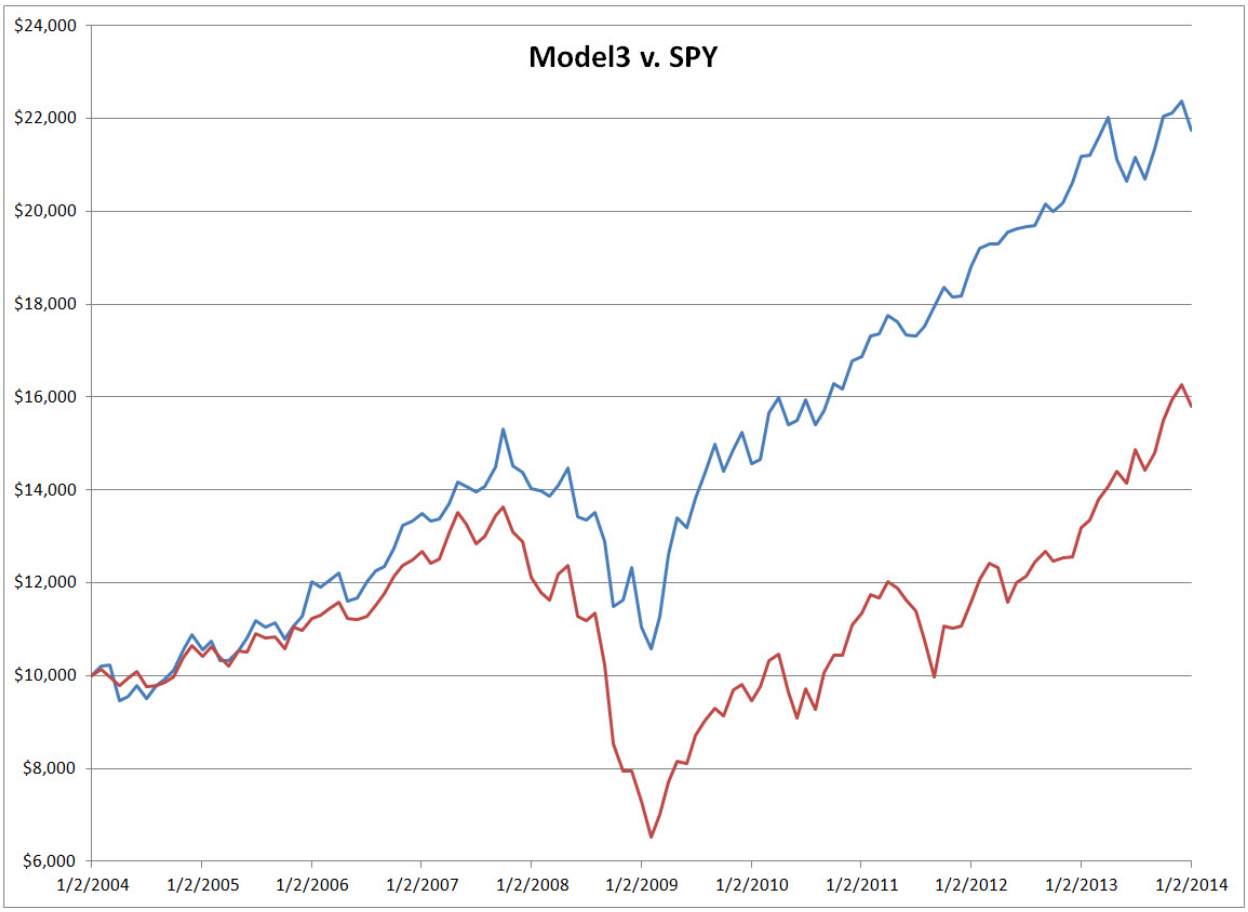

Finally, we seek a middle ground between models #2 and 3, by baselining to the *second weakest* performer. Thus, the example shown above would be baselined to -1.7%, dropping **SPY** as well as **XLB**.

SPY	EFA	IEF	TIP	EEM	IWM	XLB	IYR	TLT	Cash
0.0%	1.8%	2.7%	3.3%	2.3%	2.6%	0.0%	7.1%	2.7%	1.7%

The chart below illustrates moving the baseline. Pro-rata allocation proceeds as before, using the adjusted scores.

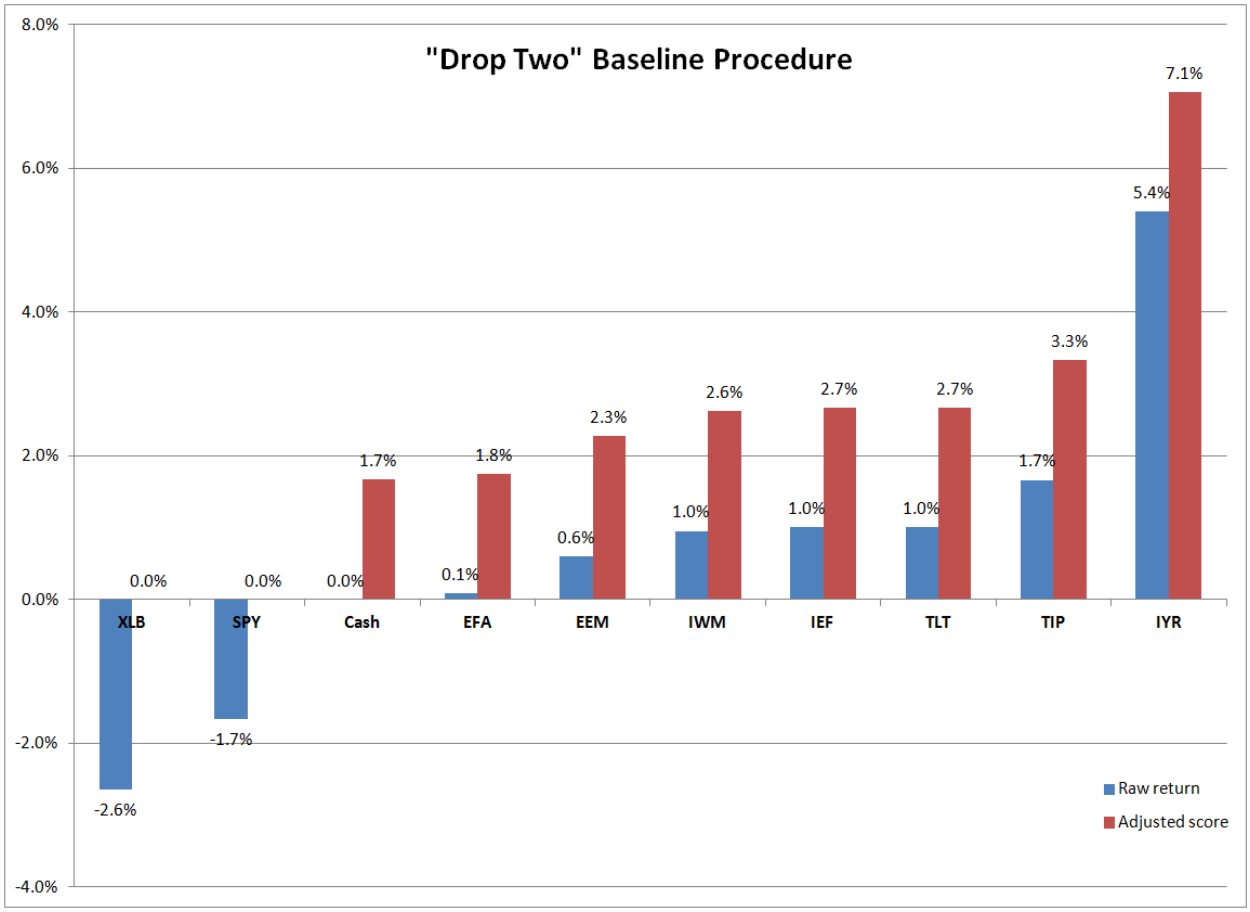

The resulting allocation is:

SPY	EFA	IEF	TIP	EEM	IWM	XLB	IYR	TLT	Cash
0.0%	7.3%	11.1%	13.9%	9.4%	10.9%	0.0%	29.4%	11.1%	6.9%

This is the "drop two weakest" model. Maintaining pro-rata exposure to eight of the ten issues (including cash) seems subjectively to be the right amount of turnover.

Below is the new allocation profile for 2005, showing the resemblance with model #2. This model tends toward the same issues, but doesn't go all in. We dispense with the exercise of dropping three, four, etc.

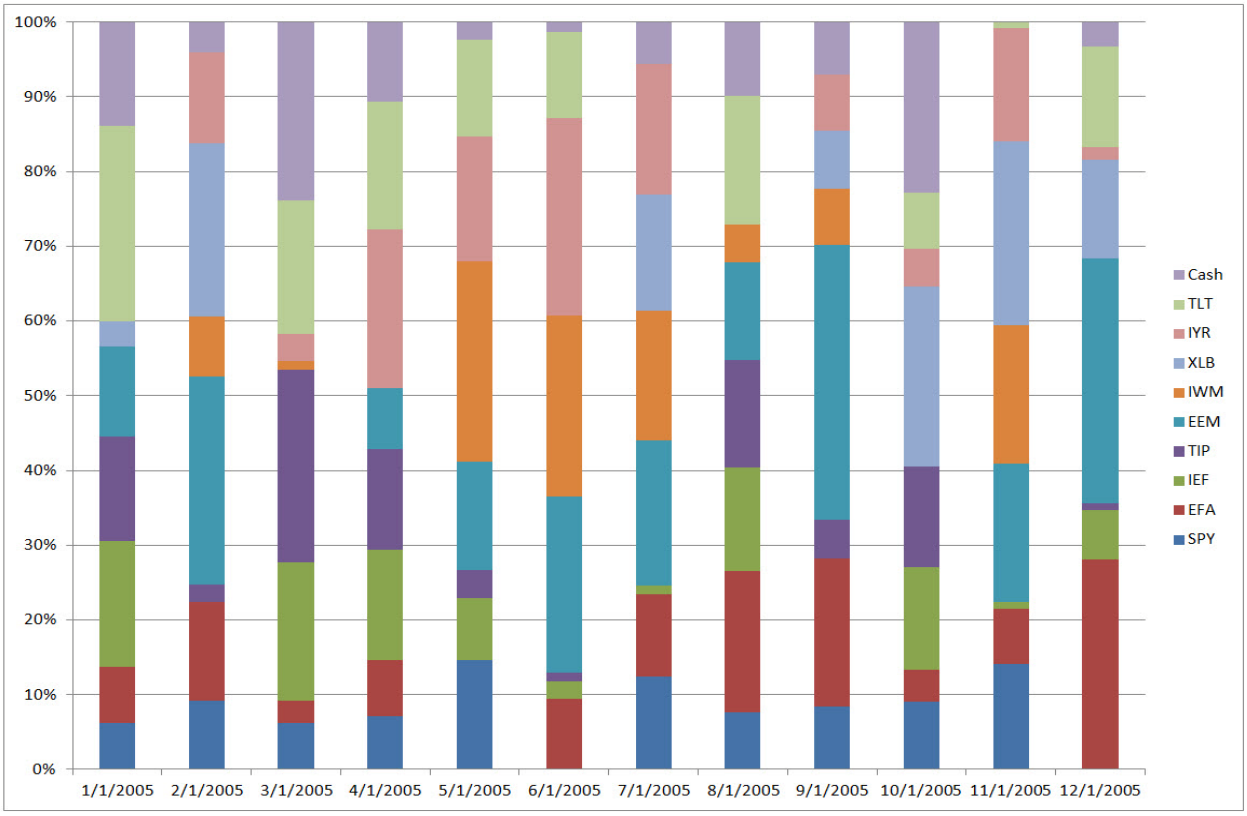

The chart below shows equity lines for all four models. Model #2 is the overall best performer, its superior return and low drawdown compensating for its higher volatility. The individual investor, however, is likely to prefer the less erratic model #4.

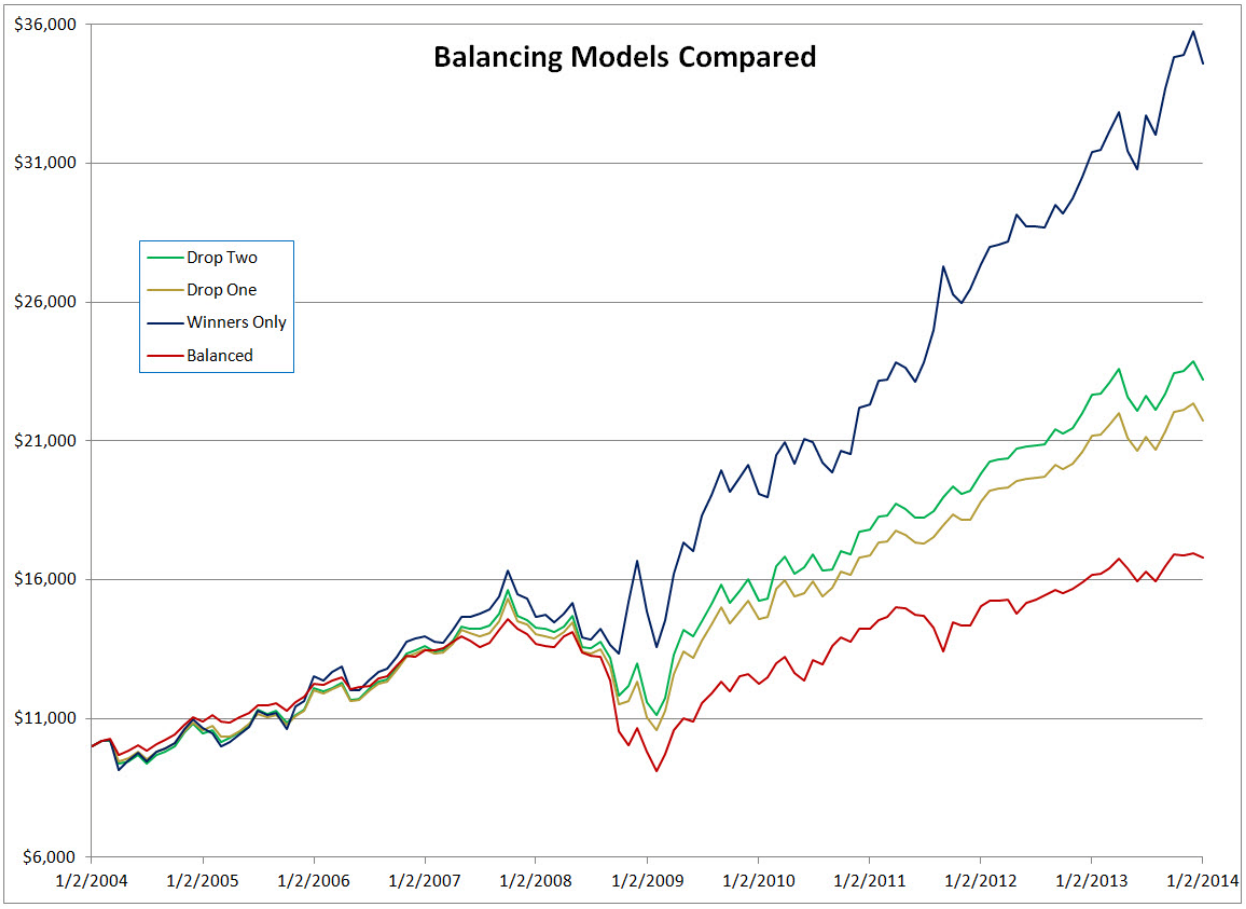

Model #1 has the smallest return and the worst drawdown. It continued allocating to stocks during the crisis of 2008. This is why many rotation models use explicit warning signals to move out of stocks. The three

pro-rata models have *implicit* warning signals that move them seamlessly into bonds and, for models #3 and 4, cash.

The table below summarizes results for the four models.

	Model1	Model2	Model3	Model4
Nickname	Equal Balance	Winners Only	Drop One	Drop Two
Total Return	67.9%	245.8%	117.4%	132.0%
Annualized Growth	5.3%	13.2%	8.1%	8.8%
Standard Deviation	11.2%	14.3%	11.4%	11.9%
Sharpe Ratio	0.40	0.86	0.63	0.67
Max Drawdown	44.7%	19.6%	35.0%	31.7%

The drawdown chart is below. The three momentum based models draw down less during the crisis, and they recoup their losses within 18-22 months versus 37 for model #1.

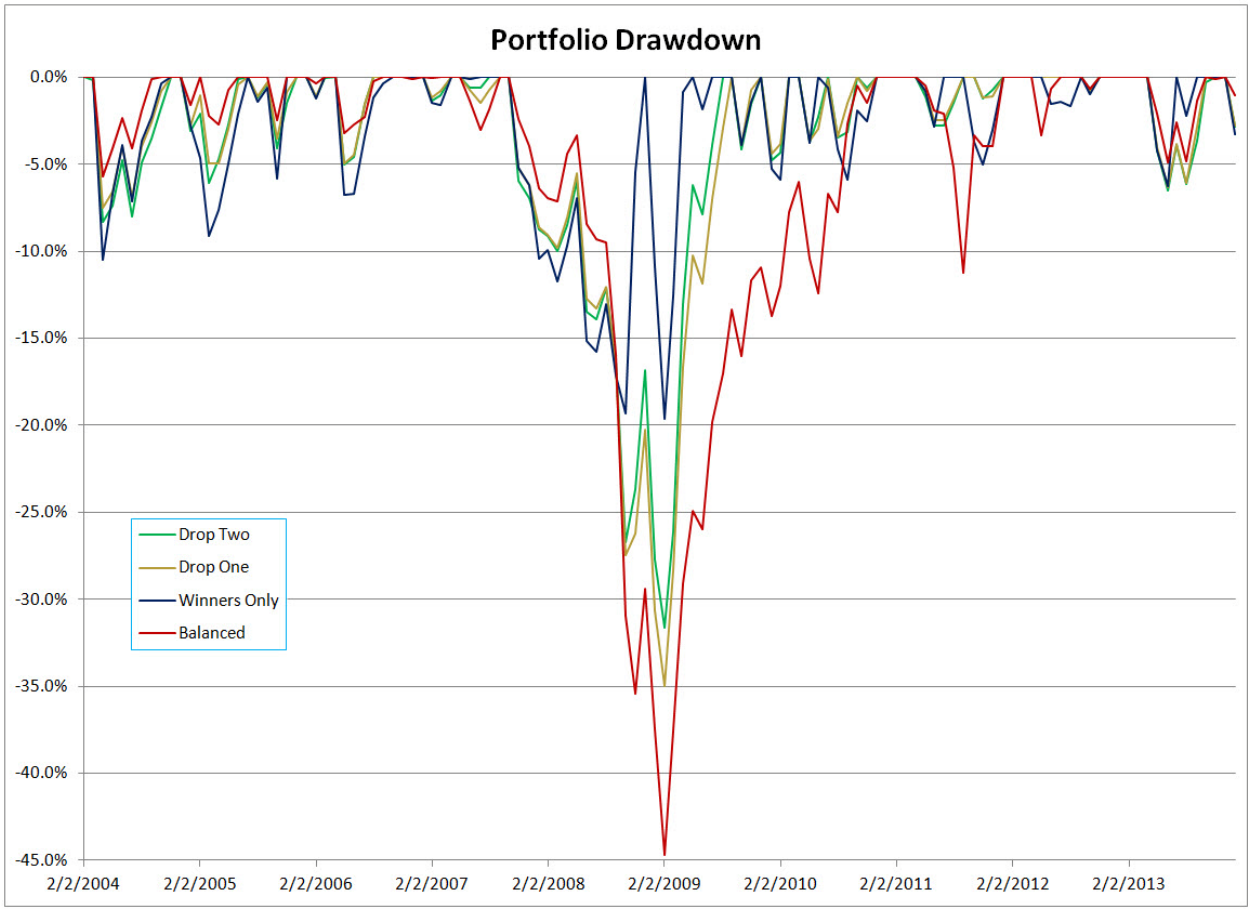

Hypothesis Test

Last month's performance is generally considered a poor guide to this month's allocation. Levy showed[3], however, that *relative rank* is predictive. Rebalancing ten issues (including cash) for 119 months results in 1,190 rank assignments, shown below with their trailing one-month returns. Each month's best performer is ranked one (1) and the worst, ten (10). Counts for each rank are not equal due to four tied rankings.

Rank	Return	Std. Dev.	Count
1	1.1%	0.052	119
2	0.7%	0.038	119
3	0.9%	0.042	120
4	0.1%	0.049	118
5	0.8%	0.046	121
6	0.4%	0.050	117
7	-0.3%	0.045	120
8	0.1%	0.053	118
9	0.4%	0.054	119
10	0.4%	0.062	119
All	**0.5%**	**0.049**	**1,190**

[3] Levy, Robert A., "Relative Strength as a Criterion for Investment Selection," Journal of Finance, Vol. 22, No. 4, (Dec. 1967), 595-610.

The relationship between rank and return is shown graphically below, with an R^2 of 0.35.

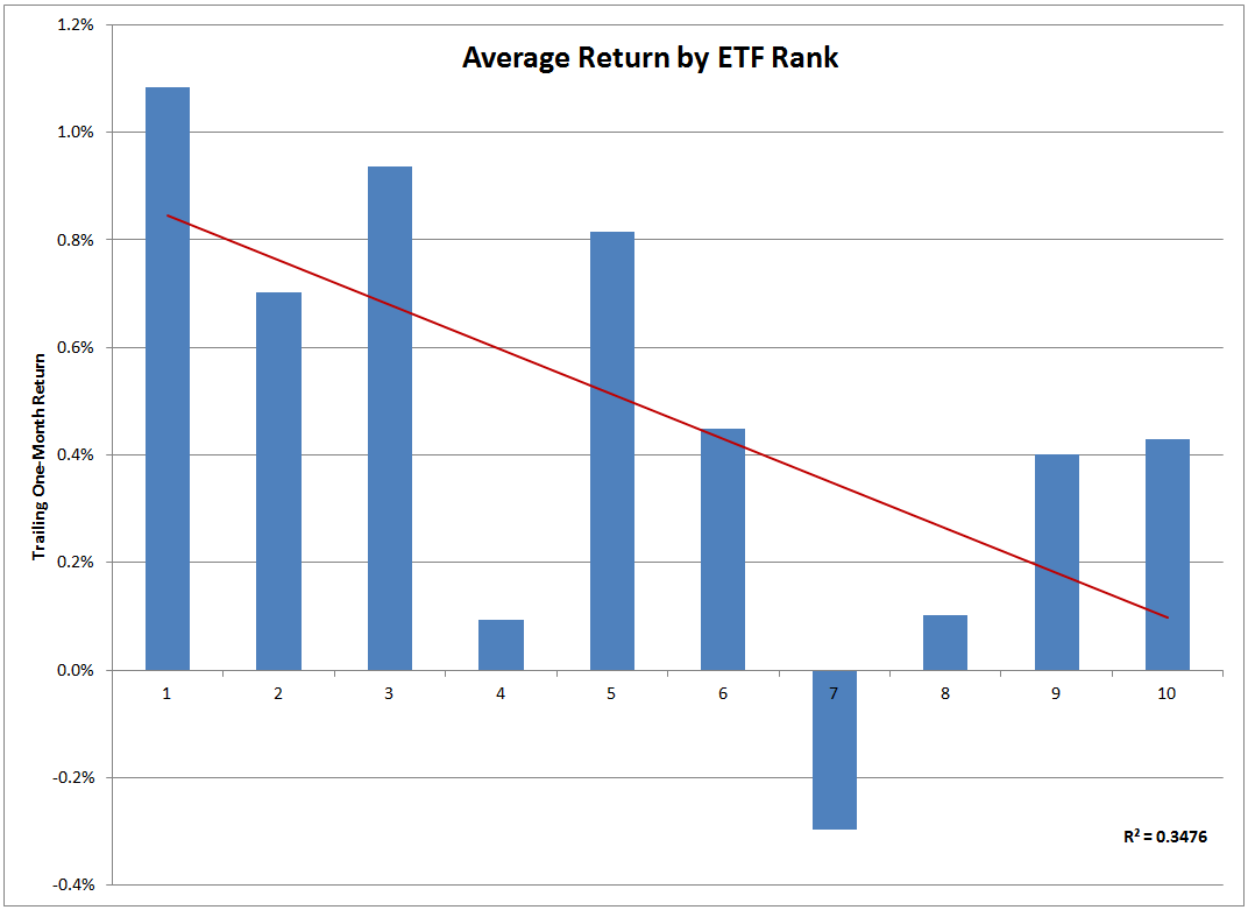

Since the three pro-rata models are generally allocating toward the better performers, and dropping one or more of the month's laggards, it seems reasonable to compare the mean trailing return of two samples:

Rank	Count	Return	Std. Dev.	z-Stat	One tail
Top Five	597	0.7%	0.046	1.786	96.3%
Bottom Five	593	0.2%	0.053		
Total	1,190				

We conclude that the higher ranked group outperforms the weaker one, with 95% confidence. Relative rank is predictive, even on a short time horizon. This brings us to the topic of inter-item correlation.

Portfolio Selection

It seems clear that the predictive effect is enhanced by weak or negative correlation among the issues. This, and the availability of data for a long term study, informed our selection of ETFs. Here are the issues we evaluated:

AGG	iShares Core Total US Bond Market	**IYM**	iShares Dow Jones US Basic Materials
DBC	PowerShares DB Commodity Index	**IYR**	iShares Dow Jones US Real Estate
EEM	iShares MSCI Emerging Markets Index	**RSP**	Guggenheim S&P 500 Equal Weight
EFA	iShares MSCI EAFE Index	**SPY**	SPDR S&P 500
GLD	SPDR Gold Shares	**TIP**	iShares Barclays TIPS Bond
IEF	iShares Barclays 7-10 Year Treasury	**TLT**	iShares Barclays 20+ Year Treas Bond
IWM	iShares Russell 2000 Index	**XLB**	Materials Select Sector SPDR

We looked at various international region funds, and settled on just **EEM** and **EFA**. We wanted to limit the portfolio to ten (10) issues including cash. Within the U.S., since we are not using sectors, we wanted a variety of capitalization sizes. Although the equal-weight **RSP** is a better choice, we chose **SPY** because it is the standard, and we compensated by including **IWM**. We tried to avoid sector funds, although we do use **IYR** and **XLB** for real estate

and commodities. This is to avoid confusion with the related topic of sector rotation[4].

For debt issues, we wanted to include government, corporate, international, and inflation protected bonds, plus various grades and durations. These features are variously covered among **AGG, TLT, TIP, BND** and **IBND**. There was insufficient history available for **EMB**.

Commodities are both economic indicators and inflation hedges. In the former category, we have **IYM, XLB, DBC, JJC**, and **USO**. Copper is correlated with the other industrial metals, of which we prefer **XLB** because it trades in higher volume. We found an international metals fund, **IRV**, but it is only five years old. So is **DBC**, a broad index of commodities. We skipped **USO** because oil prices are both a cause and an effect of inflation, and we need a pure play for the model. We tested **DBA,** to represent food price inflation.

For this study, we assume the base currency is U.S. dollars. Accordingly, we looked at currency funds, local-currency denominated bond funds, and international equity funds. We tested **GLD**, but chose **TIP** instead, as the inflation hedge. We modeled various portfolios of ten funds each. Because of

[4] Cavaglia, et al., "Risks of Sector Rotation Strategies," Journal of Portfolio Management, Summer 2001, 35-44.

the young funds, these were only five-year models. We ran the allocation models and a number of correlation studies. Here is what we found:

- The model only needs one of **XLB** and **IYM**, because they're strongly correlated. The pro-rata method will result in the same allocation going to one fund, or split between two correlated funds.
- The stock funds **SPY, IWM**, and **RSP** are strongly correlated. Using **SPY** and **IWM** is a good way to cover the gamut of capitalization sizes.
- We definitely need **EEM** and **EFA** because they're international and they have a low correlation with the other funds.
- The bond funds **AGG, TIP, IEF**, and **TLT** are correlated, plus **GLD**. As with equities, we chose **IEF** and **TLT** to cover the gamut of terms.
- We discarded **AGG** because of its low volatility. The way the model handles volatility, **AGG** is redundant with cash.
- Between **TIP** and **GLD**, we chose **TIP** as our inflation hedge. **TIP** has a longer history, and it is also convenient as the risk free rate of return in our calculations.

These findings are illustrated in the correlation matrix below, based on monthly returns from 2006 through 2013 (not the full period used in the backtest). For this period, the average of pairwise correlations among the nine chosen ETFs is 35%.

	RSP	SPY	IWM	XLB	EFA	IYM	IYR	EEM	DBC	TIP	AGG	GLD	IEF	TLT
TLT	-35%	-34%	-39%	-39%	-26%	-38%	-12%	-29%	-32%	49%	75%	12%	89%	100%
IEF	-32%	-29%	-37%	-30%	-22%	-31%	-12%	-23%	-24%	62%	80%	25%	100%	
GLD	8%	9%	6%	23%	17%	28%	8%	29%	50%	45%	33%	100%		
AGG	8%	7%	0%	2%	20%	6%	27%	16%	1%	75%	100%			
TIP	20%	20%	13%	23%	26%	27%	26%	30%	33%	100%				
DBC	54%	57%	50%	66%	61%	73%	38%	65%	100%					
EEM	83%	83%	80%	88%	91%	90%	66%	100%						
IYR	84%	79%	81%	72%	74%	70%	100%							
IYM	89%	88%	86%	97%	86%	100%								
EFA	90%	91%	84%	86%	100%									
XLB	91%	90%	88%	100%										
IWM	96%	93%	100%											
SPY	98%	100%												
RSP	100%													

Sensitivity to Lookback Period

We have presented evidence that relative rank is predictive for monthly balancing with a one month lookback. This is the most focused test of our thesis, unassisted by a longer history or historical indicators. Here, we test sensitivity to the lookback period by running model #3 with a three month lookback.

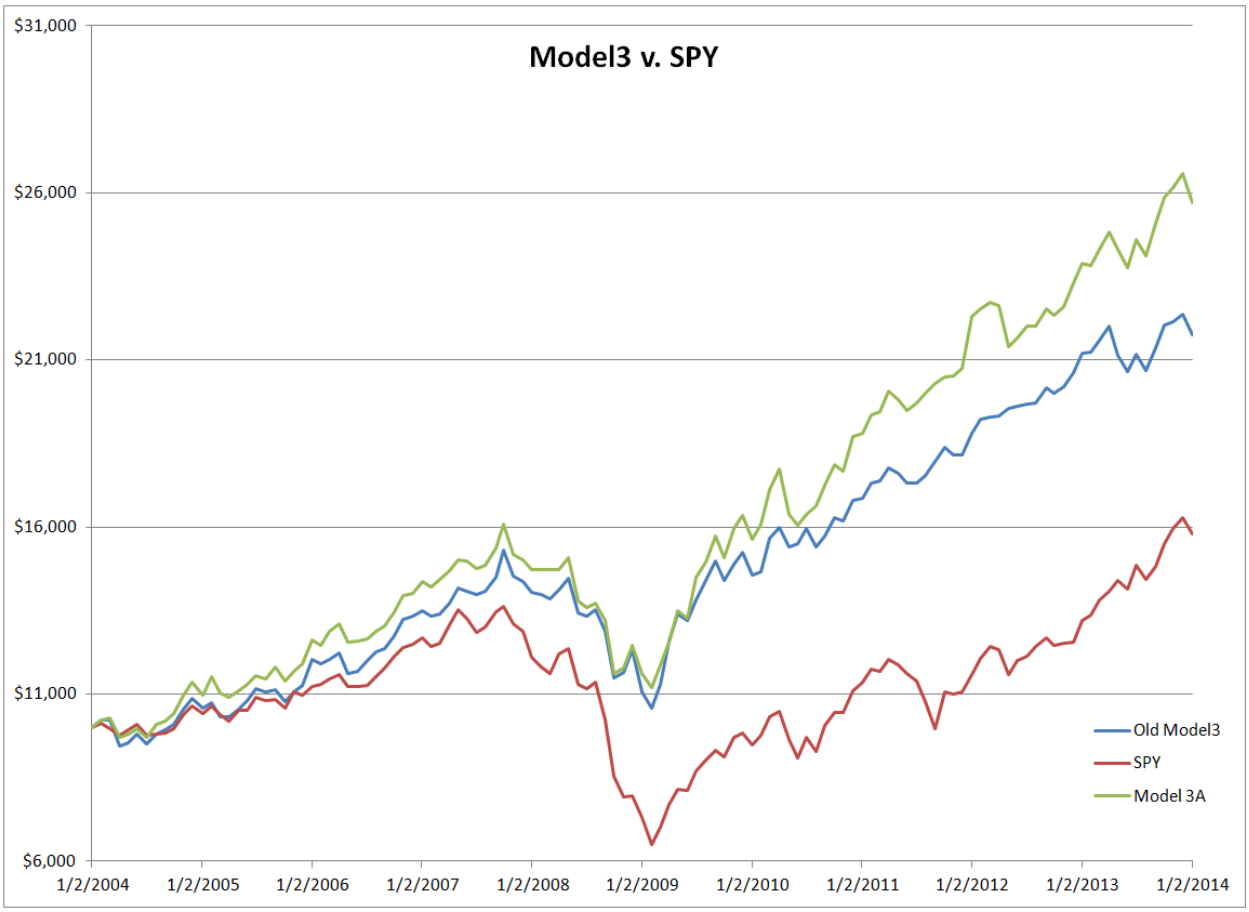

Instead of last month's return, the new model #3A allocates based on the prior three months' growth for each instrument. Lookback is not relevant for our benchmark model #1 because it always allocates in equal proportions.

The extended lookback period enhances growth with only a small increase in volatility.

	Model3	Model3A
Nickname	Drop One	Extra Lookback
Total Return	117.4%	157.0%
Annualized Growth	8.1%	9.9%
Standard Deviation	11.4%	11.9%
Sharpe Ratio	0.63	0.76
Max Drawdown	35.0%	34.2%

The new relationship between rank and return is shown graphically below, with an R^2 of 0.81.

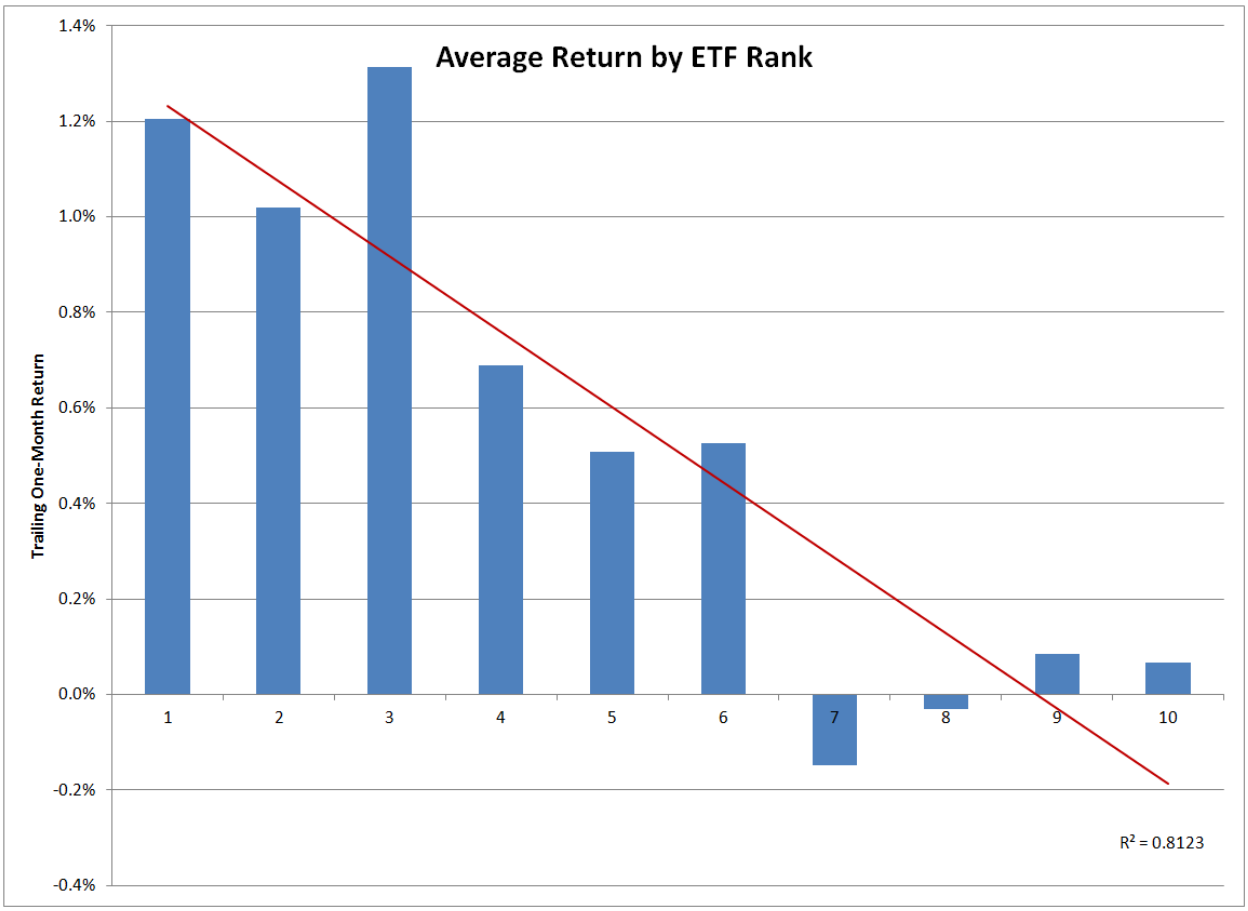

Sensitivity to Correlation

We tested sensitivity to inter-item correlation by running model #3 against the nine sector ETFs for the fifteen year period 1999 through 2013. This is a strongly correlated group, with no negatively correlated pairs. The minimum correlation coefficient is 20% (between **XLK** and **XLU**) and the inter-item mean is 55%.

	XLK	XLY	XLI	XLV	XLB	XLF	XLE	XLP	XLU
XLU	20%	34%	45%	43%	43%	40%	49%	53%	100%
XLP	26%	50%	53%	51%	46%	57%	39%	100%	
XLE	40%	45%	63%	42%	66%	47%	100%		
XLF	50%	78%	79%	64%	70%	100%			
XLB	55%	76%	85%	62%	100%				
XLV	62%	69%	69%	100%					
XLI	65%	82%	100%						
XLY	68%	100%							
XLK	100%								

SPY performs roughly as it did over the shorter period, with compound annual growth of 4.3% and volatility of 15.4%. The pro-rata model does slightly better, with growth of 5.6% and volatility of 13.8%. This is an improvement, but nothing like what we saw in the uncorrelated portfolio.

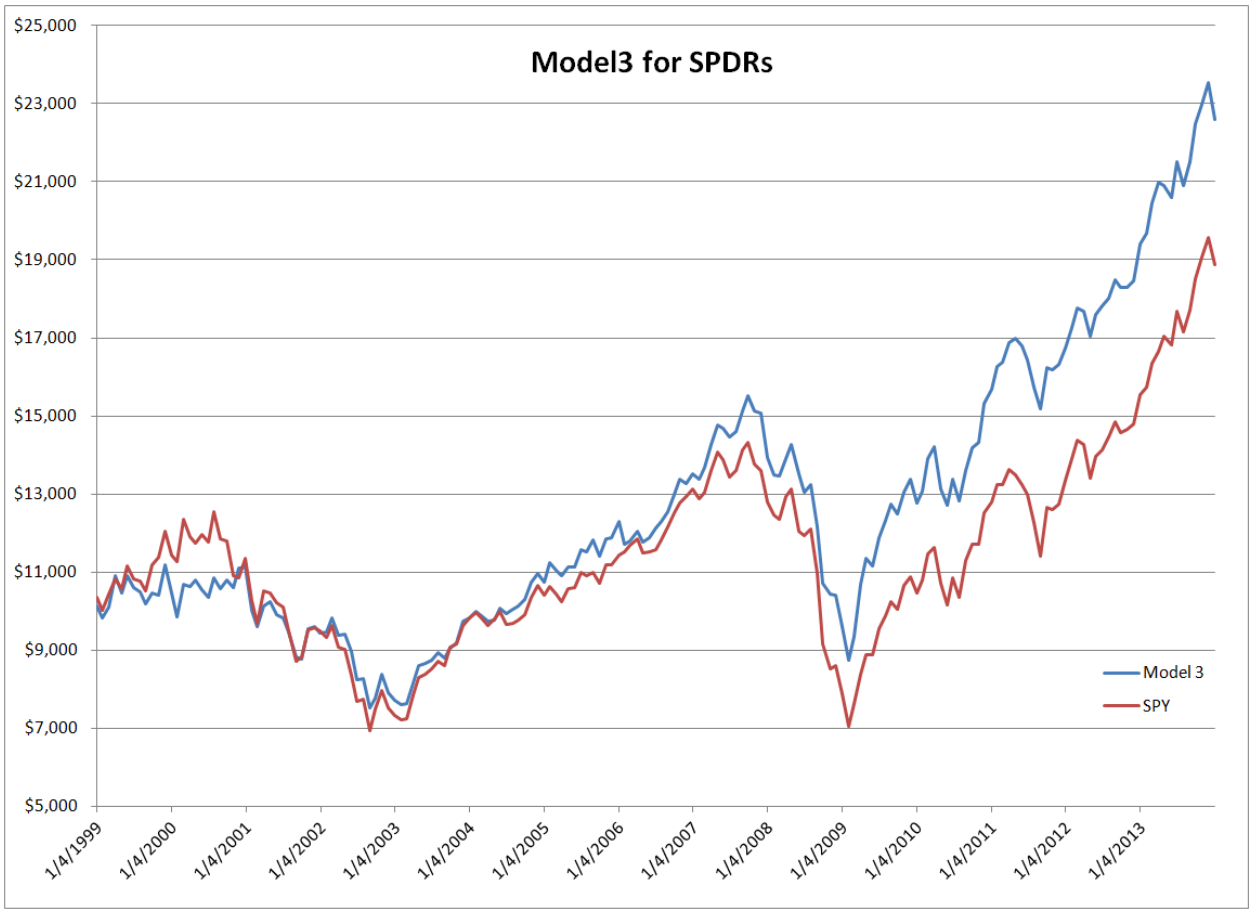

The momentum effect does not work for this portfolio because relative rank is not predictive.

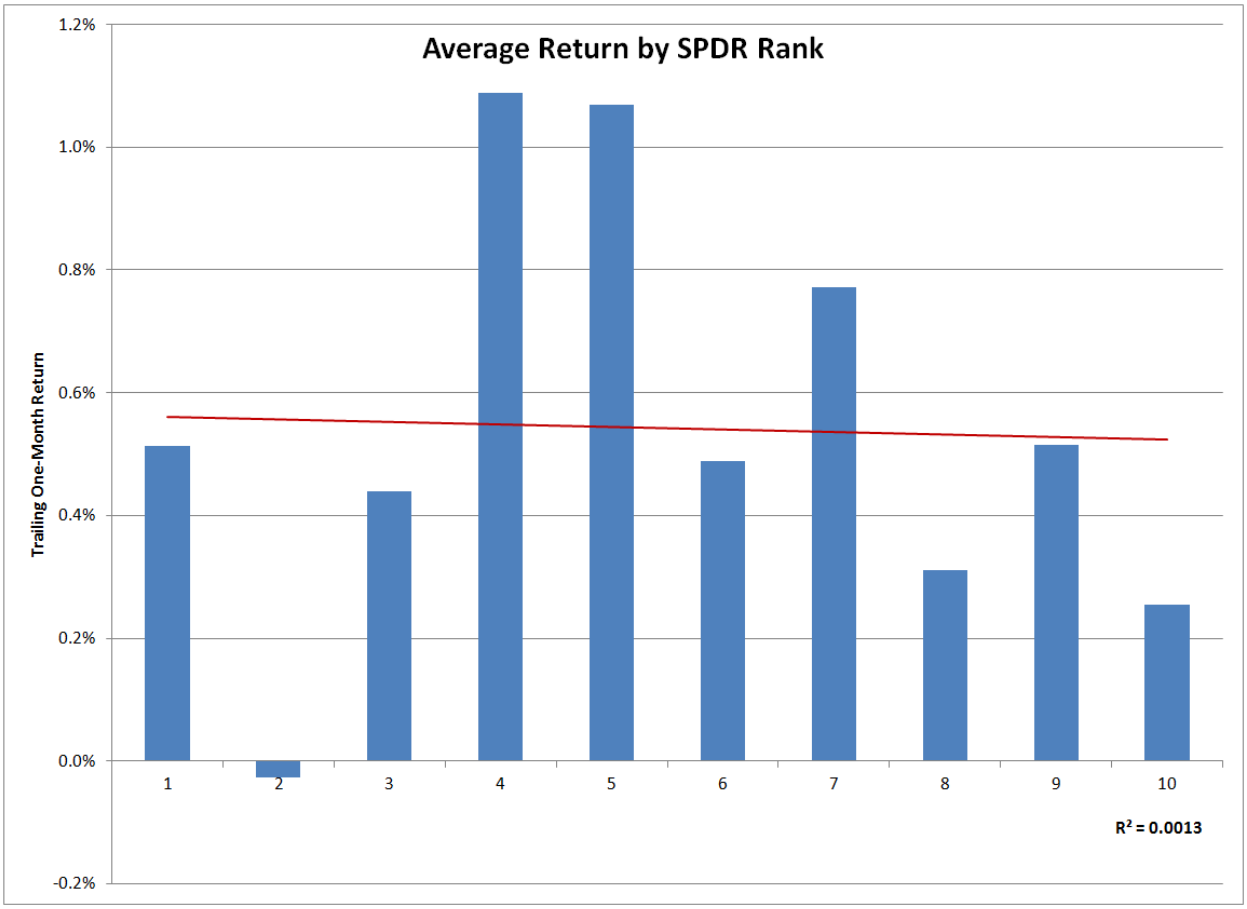

Repeating the hypothesis test confirms our suspicions. Relative rank is not predictive over the short term, where inter-item correlation is high.

Rank	Count	Return	Std. Dev.	z-Stat	One tail
Top Five	900	0.5%	0.055	-0.578	28.2%
Bottom Five	900	0.6%	0.053		
Total	1,800				

Conclusion

The three pro-rata models, using simple momentum based balancing rules, outperform the benchmark. They outperform the U.S. market, represented by **SPY**, and they outperform their individual components. We have presented statistical evidence that balancing based on short term momentum is effective because *relative rank* is predictive, but only in the context of a broadly diversified portfolio. This denotes a broad range of international asset classes, which we quantify using average inter-item correlation.

We conclude that the relative ranking theory is a sound basis for portfolio balancing, even with a short lookback period, provided that the instruments in the portfolio are sufficiently diverse. An average inter-item correlation of 35% is low enough, and 55% is not. The latter figure represents the nine SPDR sector funds, which will be familiar to most readers. For monthly balancing, a lookback period of one month is sufficient, and three months is better.

Practical Considerations

Our time series data are monthly close values from Yahoo Finance, adjusted for dividends and splits. This means that the value given for January 2004 is at month end. We assume this will also be the buy price to begin the following month, and we assume an exact dollar allocation to the ETFs, disregarding lot sizes.

The model results do not include tax and commissions. We assume that the account is a tax protected account. In an effort to reduce transaction costs, we rebalance only monthly. Model #4 stays invested in eight or nine ETFs (depending on its cash position) but it continually resizes its positions. In 120 months, it makes 1,034 total trades. Model #2, because it holds only 5.6 positions on average, makes fewer trades. Activity could be reduced by the addition of trading thresholds.

We illustrated the impact of correlation using long studies with in-sample data[5], and these may change over time. The investor cannot rely on one correlation study alone. Fundamentals must also support the identification of diverse instruments, as discussed above.

---

[5] Bandy, Howard, "Quantitative Trading Systems" (Blue Owl Press, 2007)

Implementation consists of rebalancing the nine, plus cash, every month on the second Tuesday. This avoids irregularities around month end and options expiration. Set the account preference to deposit any dividends to cash. We have organized our monthly calculator to express the target allocation as an incremental number of shares to buy or sell. It takes the current account balance as input, folding any dividends back into the model.

Procedurally, it is best to transmit all the sell orders first, raising cash, and then place the buy orders. The maximum commission for a given month is nine trades. Many brokers offer free trading in selected ETFs. For example, iShares trade free at Fidelity. It would be easy to substitute SPDR funds for their iShares counterparts, depending on the broker.