Momentum and Markowitz: a Golden Combination

Wouter J. Keller, Adam Butler, and Ilya Kipnis,

May 16, 2015 (revised June 4: v0.99d)[1]

Abstract

Mean-Variance Optimization (MVO) as introduced by Markowitz (1952) is often presented as an elegant but impractical theory. MVO “is an unstable and error-maximizing" procedure (Michaud 1989), and “is nearly always beaten by simple 1/N portfolios” (DeMiguel, 2007). And to quote Ang (2014): "Mean-variance weights perform horribly… The optimal mean-variance portfolio is a complex function of estimated means, volatilities, and correlations of asset returns. There are many parameters to estimate. Optimized mean-variance portfolios can blow up when there are tiny errors in any of these inputs. .. ".

In our opinion, MVO is a great concept, but previous studies were doomed to fail because they allowed for short-sales, and applied poorly specified estimation horizons. For example, Ang used a 60 month formation period for estimation of means and variances, while Asness (2012) clearly demonstrated that prices mean-revert at this time scale, where the best assets in the past often become the worst assets in the future.

In this paper we apply short lookback periods (maximum of 12 months) to estimate MVO parameters in order to best harvest the momentum factor. In addition, we will introduce common-sense constraints, such as long-only portfolio weights, to stabilize the optimization. We also introduce a public implementation of Markowitz’s Critical Line Algorithm (CLA) programmed in R to handle the case when the number of assets is much larger than the number of lookback periods.

We call our momentum-based, long-only MVO model **Classical Asset Allocation** (CAA) and compare its performance against the simple 1/N equal weighted portfolio using various global multi-asset universes over a century of data (Jan 1915-Dec 2014). At the risk of spoiling the ending, we demonstrate that CAA always beats the simple 1/N model by a wide margin.

1. Introduction

Imagine a high-school student, Harry, ignorant of both Markowitz and Lehman Brothers, enrolled in a class called "Finance for dummies”. At the end of august 2008 his teacher asks him to compute the so-called Efficient Frontier (return versus volatility) for a long-only mix of SPY (the SP500 ETF) and TLT (the long-term Treasury ETF) in order to determine the highest returning portfolio with target volatility (TV) of say, 10%.

Harry uses only monthly (total) return data for both assets and chooses to investigate the optimal portfolio over just the prior four months. He starts with a simple spreadsheet with only the four monthly returns (May - Aug 2008) for SPY and TLT, see fig. 1.

[1] We thank Lennart Appels, Victor DeMiguel, Frank Grossman, Winfried Hallerbach, Jan Willem Keuning , Clarence Kwan, Steve LeCompte, and Hugo van Putten for inspiration and comments. All errors are ours.

Returns	SPY	TLT
2008-05-30	1.5%	-2.7%
2008-06-30	-8.3%	2.7%
2008-07-31	-0.9%	-0.4%
2008-08-29	1.5%	2.7%

Fig. 1 Monthly returns for SPY and TLT (TR, EOM)

For each weight combination (w_{SPY}=0, 10, .., 90, 100% for SPY and the complement for w_{TLT}), Harry constructed the monthly portfolio returns, and observed the corresponding average portfolio returns ($r_P=w_{SPY}*r_{SPY}+w_{TLT}*r_{TLT}$) and volatilities of r_P over the 4 months, using simple spreadsheet functions. See fig. 2. Then he plotted both the average returns and volatilities for all 11 weight combinations as a so-called Efficient Frontier (EF). See fig. 3.

w_{SPY}	w_{TLT}	return r	volatity v
0%	**100%**	6.7%	9.1%
10%	**90%**	4.2%	7.5%
20%	**80%**	1.7%	6.3%
30%	**70%**	-0.8%	5.8%
40%	**60%**	-3.4%	6.1%
50%	**50%**	-6.0%	7.1%
60%	**40%**	-8.6%	8.5%
70%	**30%**	-11.3%	10.2%
80%	**20%**	-14.1%	12.1%
90%	**10%**	-16.8%	14.1%
100%	**0%**	-19.6%	16.2%

Fig. 2 Return and volatility for a portfolio with various SPY/TLT weights (May-Aug 2008)

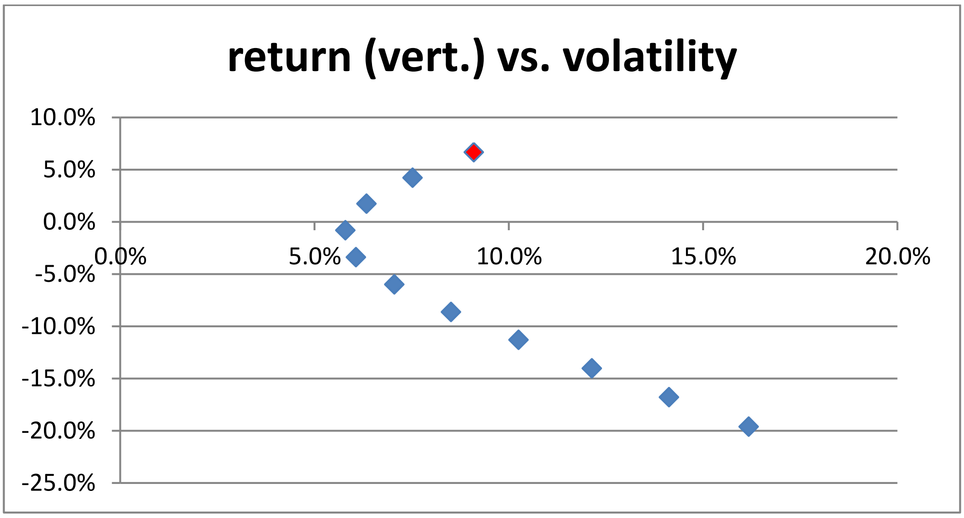

Fig. 3 The efficient frontier for August 29, 2008

As Harry is aiming at a target volatility of 10%, he chooses the weight mix represented by the far right upper point (red dot) of the EF (with an average return of 6.7% and volatility of 9.1%) for his future portfolio. This point corresponds to w_{SPY}=0% and w_{TLT}=100%. At the end of September, when Harry examines the performance of his optimal portfolio, he observes that the portfolio generated positive returns, and proved resilient to the large crashes experienced by most asset classes in that month. In fact, if he had continued to follow this process at the end of each subsequent months, he would have 'whistled past the graveyard' of the 2008 crash with only TLT in his portfolio.

Let's fast-forward 12 months from Harry's original exercise to the end of august 2009. When Harry runs the same analysis with the same lookback of 4 months (May-Aug 2009) he observes the following returns and volatilities (see fig. 4 and 5).

w_{SPY}	w_{TLT}	return r	volat v
0%	**100%**	-0.7%	8.8%
10%	**90%**	4.5%	7.6%
20%	**80%**	9.7%	6.7%
30%	**70%**	14.9%	6.0%
40%	**60%**	20.0%	5.7%
50%	**50%**	25.1%	5.9%
60%	**40%**	30.3%	6.5%
70%	**30%**	35.3%	7.5%
80%	**20%**	40.4%	8.6%
90%	**10%**	45.5%	9.9%
100%	**0%**	50.5%	11.3%

Fig. 4 Return and volatility for a portfolio with various SPY/TLT weights (May-Aug 2009)

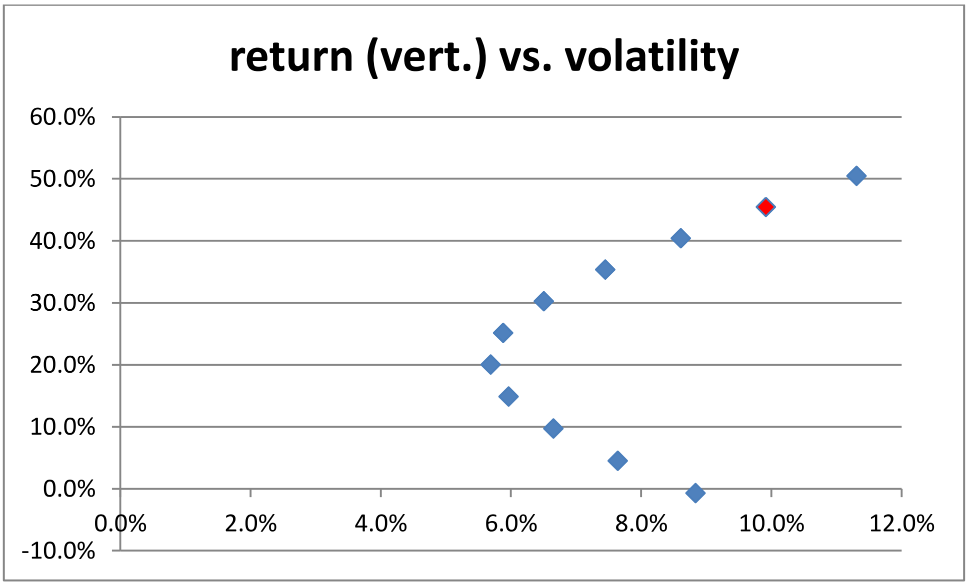

Fig. 5 The efficient frontier for August 31, 2009

Now with his target volatility of 10%, Harry chooses a portfolio of 90% SPY and 10% TLT (the red dot in fig 5.). This worked out well as Harry's portfolio benefitted from the recovery in stocks.

The positive outcomes observed from Harry's choices hints at a potential opportunity. But what has Harry done? Actually, Harry simply used Markowitz's Mean Variance Optimization (MVO), as we will show later, but with a few constraints. Specifically, he only considers portfolios with positive weights (i.e. no short-sales), and he uses a short lookback period for parameter estimation. (He also imposes discrete weights, but this is a topic for another paper).

But wait, wasn't MVO an "unstable and error-maximizing" procedure (Michaud 1989), nearly always beaten by simple 1/N portfolios (DeMiguel, 2007), in particular with limited data (here only 4 months)? And to quote Ang (2014): "Mean-variance weights perform horribly… The optimal mean-variance portfolio is a complex function of estimated means, volatilities, and correlations of asset returns. There are many parameters to estimate. Optimized mean-variance portfolios can blow up when there are tiny errors in any of these inputs. "

How can we reconcile the above statements with our simple example? It boils down to this: MVO is a great concept, but previous studies were doomed to fail because they allowed for short-sales, and applied poorly specified estimation horizons. For example, Ang used a 60-month formation period for estimation of means and variances, while Asness (2012) clearly demonstrated that prices mean-revert at this time scale, where the best assets in the past often become the worst assets in the future. In which case, it is no surprise that MVO fares poorly against simple 1/N investments, as DeMiguel (2007), Kritzman (2010), Jacobs (2013) and Ang (2014) had shown. Also, Ma (2002) shows that most of the instability of MVO disappears when we restrict ourselves to long-only.

Since the use of rolling estimation periods implies that rolling estimates should persist over the subsequent portfolio holding period, we advocate for the use of rolling estimation horizons which have historically been associated with return persistence. Of course, the phenomenon of return persistence is most consistent with the momentum factor (see e.g. Faber 2007, Hurst, 2012, Butler 2012, Antonacci 2011 and Asness 2014), which has an optimal portfolio formation period in the range of 1 to 12 months. The same holds for the persistence of volatilities and correlations (called "generalized momentum" by Keller 2012).

Fortunately, it is not even necessary to compute asset volatilities correlations, as Harry shows. It's possible to approximate the mean-variance optimal portfolio weights by permuting through discrete combinations of portfolio assets (see fig. 1) to compute the (r, v) of the resulting portfolio (see fig. 2) and arrive at the efficient frontier (fig. 3). So long as estimation horizons are in the range of 1-12 months or so, momentum will inform the return estimates. The same holds for volatility and correlation, since combining assets with opposite return patterns (i.e. negative correlation) will result in a portfolio with lower volatility (i.e. a EF point with lower v).

As Harry demonstrated with his simple spreadsheet, we don't need to bother about "complex quadratic optimizations with singular matrices" etc. In fact, as in the example above, the Markowitz model is very simple. Simply compute the efficient frontier, like Harry did, by computing the portfolio (r, v) points for various assets weights starting with the asset returns over the last few months.

2. **Calculating the Efficient Frontier: the direct and indirect approach**

Mean-Variance Optimization is at the core of Markowitz's Modern Portfolio Theory (MPT, see Markowitz, 1952). It is heavily based on the concept of the efficient frontier (EF), as explained in the above example. Even with a large universe of N assets, we can in theory examine all the possible long-only weight mixes and compute the average return (r) and volatility (v) of a portfolio over a certain short lookback period (max. 12 months). The EF is simply the envelope which encompasses all these (r, v) points. It can be shown to have the typical bullet-like shape like we see above, with the minimum volatility (MV) solution at the left-size tip of the bullet.

All solutions at the lower part of the EF bullet (below the MV point) are "inefficient" since they are all dominated (better r, lower v) by portfolios with higher returns at similar risk. So the true "efficient" frontier (EF) is the upper part of the envelope. Most "quants" focus on the Max Sharpe Ratio (MSR) portfolio which generates the maximum excess return per unit of volatility. This MSR portfolio corresponds to a point at the EF tangent to the line through the risk-free rate at the r-axis. Alternatively, one might be interested in a similar point at the EF tangent to an investor's utility function. See e.g. Ang (2014).

We prefer to focus on a constant level of risk, by using annualized *target volatility* (TV) to find the optimal point on the EF, and commensurately the right mix of weights. So an aggressive investor might choose TV=10%, while a more conservative investor might choose TV=5%. This is one of the great advantages of EF based models, like our Classical Asset Allocation (CAA). But instead of using a target volatility like we do, you can of course use the same EF (and therefore our CAA model) for a target return, for the maximization of the Sharpe Ratio, or for the minimization of the variance (or volatility), as is more common. It all depends on your preferences for risk and return.

How do we compute the EF? As said, it is simply the upper envelop of the (r, v) pairs of all possible portfolios. The computation of the portfolio return r is always very simple: just weight the various asset returns with the portfolio weights. There are two methods to compute the portfolio volatility v, given a set of weights: the direct one, where returns are generated for each asset mix (as Harry did); and the indirect one, using a covariance matrix. This latter method, which is employed by most quadratic optimizers, requires the calculation of the asset covariance matrix C over lookback horizon T. It can easily be shown[2] that when C is estimated over the same lookback period, the indirect method using C is equivalent to the direct method Harry used in his spreadsheet.

When N (the number of assets) is much larger than T (the number of observations), so N>>T, it is computationally more efficient to use Harry's method for the portfolio volatility (given a set of weights); when N <<T it is better to compute the volatility on the basis of the covariance matrix. The

[2] In matrix notation: If C is the historical NXN covariance matrix for N assets, X the TxN (demeaned) matrix of returns with T observations (T=12 when using the last 12 months), and w the N-vector of weights of a portfolio, the variance of that portfolio equals v^2= w'X'Xw/T= w'Cw since C=X'X/T. So the historical variance (and therefore the volatility v) of the portfolio with weights w can be computed directly by using the historical variance of the portfolio over the last 12 months (i.e. the variance of Xw) or indirectly by first computing the covariance matrix C and from there the variance v^2=w'Cw/T. This also shows that the variance (and volatility) is a quadratic function (w'Cw) of the weights w.

results are the same. Notice that with our monthly momentum model (T<=12), Harry's method is more efficient for all universes larger than N=12. It is also much simpler to understand, since there is no covariance matrix involved, just the historical volatility of the portfolio over the last T<=12 months.

Notice that Harry's method works great even when the number of assets N is much larger than T: simply generate random assets weights (long-only), compute the portfolios return r and volatility v over e.g. the last T observations and plot these (r,v) pairs for all generated weights in the EF diagram. There are no "complex function of estimated means, volatilities, and correlations of asset returns" and not "many parameters to estimate" (Ang, 2014).

As Harry's example shows, the computation (using the direct method) of the EF could not be simpler. And as we will show, the combination of a momentum based tactical approach (T<=12 months lookback) with a long-only constraint transforms Markowitz' MVO from an "unstable and error maximizing procedure" into a robust portfolio formation method.

In addition, we can easily show that both the EF and the corresponding optimal weights for a given target volatility are invariant for changes of scale (say monthly or annual) in returns and volatilities while they are also independent of parallel changes in the level of returns for all assets (which simply moves the EF up or down)[3].

So, the optimal weights (corresponding to the point on the EF with a given target volatility) are simple to understand and often very robust, in particular for long-only portfolios. This can be demonstrated with Excel. Even with a universe with N=9 assets and a short lookback of T=4 months as Harry used (so T<<N) you can simply generate several thousand random (non-negative) weights in Excel to see the EF as the upper envelop of the resulting scatter. Just compute the portfolio (r, v) for each set of weights, given the 4x9 returns. Now you will also see that there is not much danger that "Optimized mean-variance portfolios can blow up when there are tiny errors in any of these inputs" (Ang, 2014). Even when N>>T, the EF (and the optimal weights for a given TV) are very robust under our assumptions (T<=12 months and long-only). Of course, for different points in time and for different universes the optimal weights will be different. But as you will see, you will always get the well-known EF bullet shape.

Actually, this direct EF approach using a spreadsheet with random weights is simple to understand and rather straightforward, and shows the beauty of "tactical, long-only" MVO. However it can be computationally very, very slow, especially if N is large and you have to compute the EF (and the optimal weights for a given TV) for each month walking forward through your backtest. That's why more experienced quants use numerical programming methods.

Analytically, given a target volatility, finding the optimal weight mix is a quadratic optimization problem under constraints, since the volatility is a quadratic function of the weights (as we explained in footnote 2). Traditionally, an investor might use a complex quadratic optimizer (with constraints)

[3] So you will find the same optimal portfolio weights with CAA and a target volatility when using excess returns instead of normal returns (given a riskfree rate per month). Notice that the MSR solution is, however, not invariant for these changes unless the level of the riskfree rate is also changed by the same amount as all the returns. And of course this invariance only holds when all the returns are shifted by the same amount for a given month. See also par. 4 below.

to find the optimal portfolio solution (like the one corresponding to the target volatility on the EF). However, many quadratic optimizers have problems with the combination of a short estimation history (e.g. our 12 months) with many assets (much more than 12). There is, however, a better alternative: the *Critical Line Algorithm (CLA)*.

3. Markowitz's Critical Line Algorithm

Markowitz's Critical Line Algorithm (CLA) is one of the most elegant and fastest algorithms to calculate the efficient frontier under constraints (see Markowitz 1956 and 1959, Kwan 2007, and Niedermayer, 2007). It is also effective when the number of assets N is much larger than the number of observations T (see Nawrocki, 1996), in contrast to many (traditional) quadratic MVI optimizers.

The main idea behind the CLA algorithm is very simple: you start at the upper right side of the EF with the asset with the highest return, and follow the EF to the left (towards the minimum variance MV solution) by searching for the best asset to be added (IN) or subtracted (OUT) from the portfolio one at the time. The EF then consists of many curved lines (called the "critical lines") connecting the so called "corner points" where assets change from IN to OUT of the portfolio or vice versa. Over the critical lines, and between the corner points, the assets included in the portfolio don't change, only their weights. The collection of all critical lines (and corner points) builds up the EF, starting from the upper right point to the MV solution at the far left. Notice that, although the critical lines between two corner points are linear on weights (much like fig. 2), they are curved (parabolic) in the (r,v) space of the EF (as shown in fig. 3). In Appendix A we show some examples of the EF produced by CLA for a given month and universe with the various corner portfolios.

Many traditional MVO optimizers have problems in cases where the NxN covariance matrix is "singular"[4]. This singularity of the covariance matrix arises when the number of assets N is larger than the number of observations T. Our use of tactical (short-term) instead of strategic (long-term) lookbacks means that T is limited to 12 (with monthly observations), in order to take price momentum into account. This might present issues for traditional quadratic optimizers, but not for CLA.

For CLA, as long as the number n (<=N) of assets in the optimal (TV) portfolio at a critical line segment are less than the number of observations T (so n<T), there exists a unique solution for CLA. And since there was no open source version of CLA available in R, we (Ilya) coded it ourselves (see Appendix B) and made it available to the public, including the code for our CAA model.

Because of the stepwise approach over all the corner points in CLA, only a linear (instead of a quadratic) problem in a small dimension (n<<N) has to be solved for each corner point, where n is the number of assets in the portfolio around the corner point. See Kwan (2007) for an illustration of CLA in Excel, where also the handling of simple constraints on the weights are shown, such as no short-sales ("non-negative weights") and investment limits ("weight caps"). [5] Using CLA, we had no

[4] A covariance matrix C is singular if the inverse C^{-1} does not exist. This is the case when T<N.

[5] Our simple open-source CLA algorithm in R (implemented by Ilya Kipnis), is based on this article of Kwan (2007). Our algorithm is the first open-source implementation of CLA (and CAA) in R as far as we know. More optimized (open-source) CLA implementations can be found in Niedermayer, 2007 (in Fortran) and Bailey, 2013 (in Python).

trouble handling large universes (e.g. N=39, shown below) notwithstanding our limited number of (monthly) observations (T=12).

4. **Properties of CAA**

Our Classical Asset Allocation (CAA) model applies mean-variance optimization over a momentum-friendly lookback horizon of less than one year. We have already discussed the invariance of the optimal allocations of the MVO model (and therefore of our CAA model) when we change the *scale* of measurements of returns and volatilities from yearly to say monthly. The optimal CAA allocation for a target volatility is also independent of the absolute *level* of returns. Subtracting e.g. the T-Bill rate in a month from all other assets returns will not change the outcome: the EF is simply moved up or down when the level of all returns is changed by the same constant and the optimal allocation for a given target volatility stays unchanged[6]. Both characteristics (scale and level independence) contribute to the robustness of CAA. (We should note that the max Sharpe portfolio is not scale invariant, as the point of tangency of the EF with the Capital Allocation Line will change as the EF moves higher or lower on the risk-return plane.)

An important property of any asset allocation model is *independence of irrelevant alternatives* (IIA): we should have the same optimal portfolio result (in terms of r, v) when linear combinations of its existing universe assets are added to the portfolio. So when (a copy of) SPY is added to the universe which already holds SPY, the result should not change.[7] See also Choueifaty, 2011. With MVO (and therefore with ditto CLA), any mix of two perfectly correlated assets (with correlation one) leads to the same EF and the same portfolio result, so any model based on MVI obeys the IIA principle, and so does our CAA model.

Although we focus with our CAA/CLA implementation on the optimal allocation with a given target volatility (TV), we might also choose the point on the EF corresponding to the Maximization Sharpe Ratio (MSR), given a risk free rate, as is often done with MVO. We will call this the MSR variant of the CAA model. We then arrive at most well-known *"smart-beta"* allocations as special *submodels of* this MSR variant of CAA. These smart beta models include Equal Weight (or 1/N), Minimum Variance (MV) or minimum volatility, full Risk Parity (full RP, aka ERC) and naïve Risk Parity (naïve RP), and Maximum Diversification (MD, see also Choueifaty, 2011). These are all special cases of the general Maximum Sharpe Ratio (MSR) solution of MVO, given certain assumptions for the returns and the covariance matrix, and therefore also special cases of CAA/MSR. These submodels are in particular of interest when we are not willing to specify expected returns in detail. Notice that this is less relevant when momentum is used, as in our implementation of CAA below. All these smart beta models are independent of the risk free rate, here assumed to be zero.

As is shown in Hallerbach (2013), the EW solution corresponds to the CAA/MSR solution when we assume all (expected) returns, volatilities, and correlations to be equal over all assets. The MV solution corresponds to the CAA/MSR solution when all asset returns ri are assumed to be equal. Because of the level independence of the returns, running CLA with ri=r for any r will give the MV

[6] We thank Clarence Kwan for this insight.

[7] It is clear that our previous models (see Keller 2012, 2013, 2014) do not possess this desirable property since adding SPY will simply duplicate SPY if it is part of the optimal portfolio, i.e. give it a double weight.

solution, given the historical covariance matrix. This is also (of course) the left point of our EF, as computed by CLA. We arrive at the MD solution as special case of CAA/MSR when all asset returns are proportional to their individual volatilities (MD), implying equal Sharpe ratios over assets. Were we to replace ri with volatility vi per asset and run CLA/MSR, we would arrive at the MD portfolio.

If, in addition, all cross-correlations are assumed to be the same and equal to the average historical correlation over all assets, we arrive at full RP (or ERC, see e.g. Maillard 2008), while when assuming zero correlations (still in addition to constant Sharp ratio's) we arrive at naïve RP. So several smart-beta models can be found as special cases of CAA/MSR, in particular when we are not willing to use momentum for returns. For the rest of this paper, however, we will use momentum with CAA as an added advantage to MVO.

5. Data

Total return index histories for 39 asset (classes) were assembled from several sources. We refer to Keller (2014b) for details. Returns were drawn from Global Financial Data (GFD), Kenneth French Database (FF), Barclays, MSCI, Yahoo and other providers of historical data.

The assets (classes) have been divided into three overlapping universes of increasing size:

N=8 universe:

- S&P 500
- EAFE
- Emerging Markets
- US Technology Sector
- Japanese Equities
- 10-Year Treasuries
- T-Bills
- High Yield Bonds

N=16 universe:

- Non-Durables
- Durables
- Manufacturing
- Energy
- Technology
- Telecom
- Shops
- Health
- Utilities
- Other
- 10-Year Treasuries
- 30-Year Treasuries
- U.S. Municipal Bonds
- U.S. Corporate Bonds
- U.S High Yield Bonds
- T-Bills

N=39 universe:

- S&P 500
- US Small Caps
- EAFE
- Emerging Markets
- Japanese Equities
- Non-Durables
- Durables
- Manufacturing
- Energy
- Technology
- Telecom
- FTSE Global ex-US
- FTSE Developed Equities
- FTSE Emerging Markets
- 10-Year Treasuries
- 30-Year Treasuries
- U.S. TIPs
- U.S. Municipal Bonds
- U.S. Corporate Bonds
- U.S High Yield Bonds
- T-Bills
- Int'l Gov't Bonds

- Shops
- Health
- Utilities
- Other Sector
- Dow Utilities
- Dow Transports
- Dow Industrials
- FTSE US 1000
- FTSE US 1500
- Japan 10-Year Gov't Bond
- Commodities (GSCI)
- Gold
- REITs
- Mortgage REITs
- FX (1x)
- FX (2x)
- Timber

6. Implementing the Classical Asset Allocation (CAA) model

Now we will specify the implementation of our *Classical Asset Allocation* (CAA) model in more detail. It is, as will be clear by now, based on Markowitz MVO, with a short lookback period (maximum 12 months) and a no-short-sales constraint (long-only). We will use CLA to compute the optimal allocation on the EF, given a target volatility (TV).

Each month, we estimate the optimal mix of assets weights based on (tactical) information from the prior 1 to 12 months (to have momentum effects) and use that mix for the next month. As a result, since our simulations use a century of data, we performed 100 years * 12 rebalances per year = 1200 optimizations. This would have been computationally intractable with larger universes using Harry's method in Excel, but with our (not very optimized) CLA code the process took just a few seconds.

For each optimization we take the rolling historical means (returns) and covariances as proxies for the expected returns and covariances next month. Faber (2010), Butler (2012), Antonacci (2011, 2013) and others suggest that, in fact, the momentum effect is robust over a range of approximately 1 month through 12 months. As a result, we chose to average the momentum observed over 1, 3, 6, and 12 months in order to span the optimal space and stabilize our estimates. This is also consistent with the lookback specification for EAA (Keller, 2014b).

For the covariance matrix (i.e. volatilities and correlations) we used the historical covariance matrix of returns for the trailing twelve months. We also tried using daily instead of monthly data for the covariance matrix but found no significant advantage[8]. This is a topic for further research. Notice that all returns are total returns, i.e. including dividends etc.

As our implementation of mean-variance optimization seeks any optimal set of portfolio weights, there is the potential for the portfolio to become quite concentrated at times. To reduce this possibility, we imposed caps (max weights) on all risky assets to enforce greater ex-post diversification. E.g. with a universal cap of 25% or 50% for all assets, the portfolio should contain at least resp. four or two assets (i.e. with non-zero weights). This forced diversification also often improves the ex-post Sharp-ratio (see Appendix A). Our default will be to impose a cap of 25% for all risky assets and no cap (i.e. a cap of 100%) for all cash-like assets. As a result the CAA model can

[8] Daily data might prevent singularity of the covariance matrix (much larger T) but even for our largest universe (N=39) this was not a problem with our monthly data and CLA. Also clustering (autocorrelation) will be more of problem with daily data. And for a century of data (incl. rate hikes etc.) we only have monthly data at our disposal.

move fully into “safer” assets at turbulent times. We will use both 3-month T-Bills and USGov10y bond as cash, so both are uncapped in all simulations.

All simulations are performed using monthly data from Jan 1915 through Dec 2014 (with one year lookback, i.e. from Jan 1914), so over a full century, with monthly rebalancing. We will always use the EW (or 1/N) allocation as our naïve benchmark. We will report per backtest:

- R (CAGR),
- D (maximum Drawdown at monthly observation frequency over a century),
- V (annualized Volatility),
- SR5 (Sharpe Ratio with 5% threshold, so SR5=(R-5%)/V)
- CR5 (Calmar Ratio with 5% threshold, so CR5=(R-5%)/abs(D)),

The target rate of 5% for the Sharpe and Calmar Ratio approximate the average 3m T-Bill rate over the 100 year test horizon. Notice that the number of free parameters (degrees of freedom) to influence the backtest performance are minimal, especially when compared with other tactical models, which often require individual parameter weightings (see Keller 2014b); pre-specified number of holdings or binary ‘in or out’ thresholds, and; weighting schemes for portfolio assets. In contrast, our tests involve just four assumptions, as the optimization selects for both number of assets and weights:

- Long only, so no short sales
- The estimation periods for monthly returns: equal for 1,3,6, and 12 months, zero else
- The estimation period for the monthly covariances: equal for the last 12 months
- The size and asset compositions of the testing universes: N=8, 16, and 39

We would note that the long-only constraint was imposed to make the implementation practical for most readers, and to improve the robustness of the solution, see Ma (2002). The momentum lookback parameters were chosen to span the known effective momentum range and to minimize the potential for curve-fitting. The covariance estimation horizon was chosen to correspond to our longest lookback period in order to maximize the density of the covariance matrix. Lastly, our goal in testing three diverse universes is to avoid data-snooping an optimal universe. Overall, our goal was to minimize the number of free parameters to broadly minimize the potential for curve-fit results, and maximize the likelihood of persistent performance out of sample.

To this small list we will add two simple parameters: one max weight (or “cap”) for all ‘risky’ assets, and the target volatility (TV) to arrive at the offensive or defensive model.

- A maximum weight (“cap”) of 25%
- The Target Volatility: 10% (for offensive) and 5% (for defensive models)

We assume that all assets are “risky” except T-Bills and USGov10yr, both of which are considered as non-risky or “cash” and are therefore uncapped (or have cap 100%). All other risky assets will have a default cap of 25% in order to produce palatable portfolios and enforce a minimum amount of diversification. Notice that a cap of 25% for risky assets implies that we will have at least 4 risky assets in the portfolio when there is no cash selected and more or less assets when (uncapped) cash is selected.

The target volatilities are more model variants than optimization parameters. In addition, we show in Appendix A that our results (in terms of dominance over the benchmark EW=1/N) are very robust for other choices of cap=25%. As there are effectively no degrees of freedom left for datasnooping, in particular when compared to all the free parameters (wR, wV, wC, wS) in our previous models (see Keller 2012-2014), we will not present an in-sample (IS) versus out-of-sample (OS) comparison as we did before (see Keller, 2014b). So the whole century (Jan 1915-Dec 2014) will be used for a single extended simulation.

We will now present the results for the three universes: a small (N=8) global multi-asset universe, a medium size (N=16) universe with US sectors and bonds, and a large (N=39) global multi-asset universe.

7. The small (N=8) universe

Here we will present the results for CAA for the small (N=8) universe for Jan 1914-Dec 2014 (100 years). This universe consists of: SP500, EAFE, EEM, US Tech, Japan Topix, T-Bills, US Gov10y, and US High Yield. All non-cash assets will have the default 25% cap, while both “cash” assets T-Bills and US Gov10y are uncapped. The benchmark is the equal weight (EW or 1/N) model of the full portfolio

The numerical results for the offensive (TV=10%) and the defensive model (TV=5%) models are shown if fig. 6 and 7, resp. in the graphs in fig. 8 and 9.

TV=10% (N=8)	CAA	EW (1/N)
R=	12.7%	8.7%
V=	8.3%	9.2%
D=	-17.3%	-49.7%
SR5 (Sharpe)	92.3%	40.1%
CR5 (Calmar)	44.6%	7.4%

Fig. 6 The offensive (TV=10%) results for N=8 (Jan 1915 – Dec 2014

TV=5% (N=8)	CAA	EW (1/N)
R=	10.5%	8.7%
V=	5.8%	9.2%
D=	-10.6%	-49.7%
SR5= (Sharpe)	94.5%	40.1%
CR5= (Calmar)	51.5%	7.4%

Fig. 7 The defensive (TV=5%) results for N=8 (Jan 1915 – Dec 2014)

As can be seen in both tables (fig. 6 and 7) and graphs (fig. 8 and 9), the CAA performance is excellent, in particular taking into account the lack of free parameters, except for the default cap of 25%, which is rather conservative (see Appendix A for more aggressive alternatives like cap=100%).

For both the offensive (TV=10%) as the defensive (TV=5%) model, the returns R, volatilities V, max drawdown D , Sharpes SR5 and Calmars CR5 are all substantial better than the (EW=1/N) benchmark. For the offensive model the return R is nearly 50% improved over EW (12.7% vs. 8.7%), while volatility V is slightly less and in particular maximum drawdown D is reduced by almost two thirds (-17% instead of -50% for EW). The largest CAA drawdowns were on Black Monday (Oct 1987) and at the start of WW2 (May 1940), see fig. 9. Even for the defensive model the return R is better than EW (10.5% vs. 8.7%), while volatility (5.8 vs 9.2%) and drawdowns (-10% vs -50%) are half and resp. one fifth that of the benchmark.

For return/risk statistics, both models (offensive TV=10 and defensive TV=5%) have Sharpe ratio’s (SR5) more than double that of the EW, with both optimizations yielding Sharpe ratios above 0.9 vs. 0.4 for EW. The Calmar ratio (CR5) improves more than five times over the benchmark (CR5=0.07 for EW).

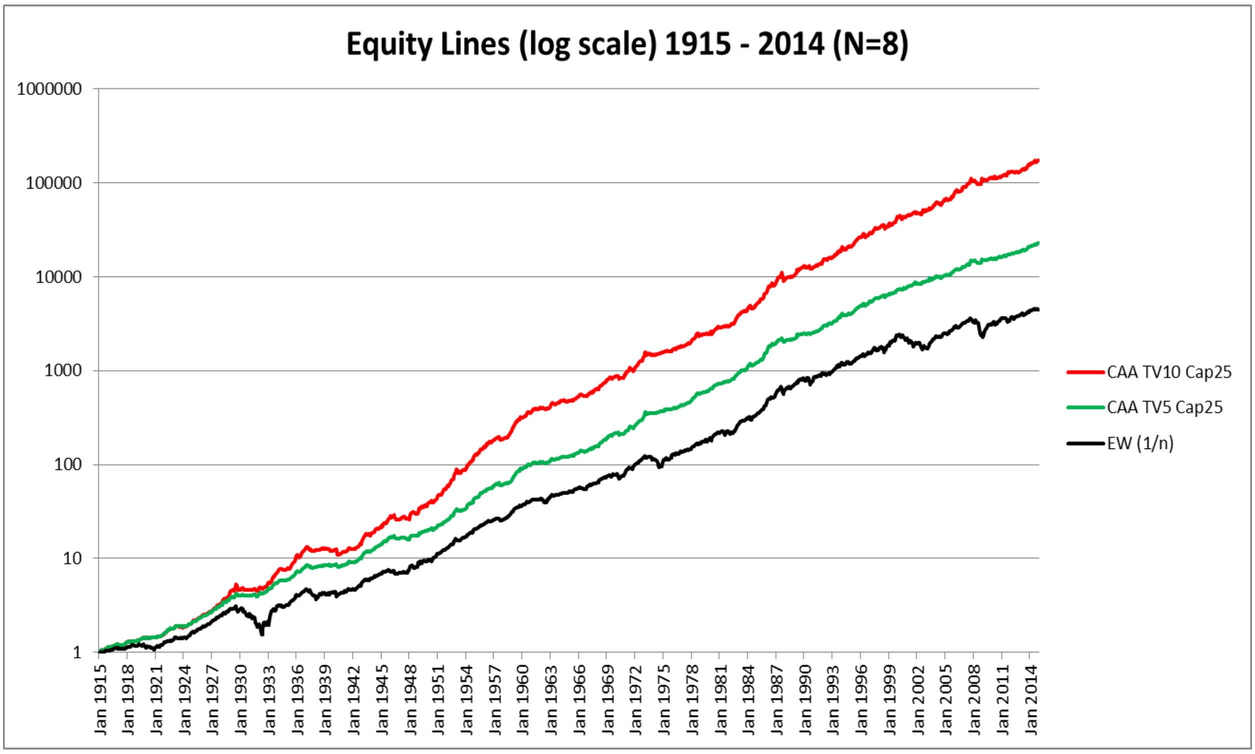

Fig. 8 Equity lines for N=8

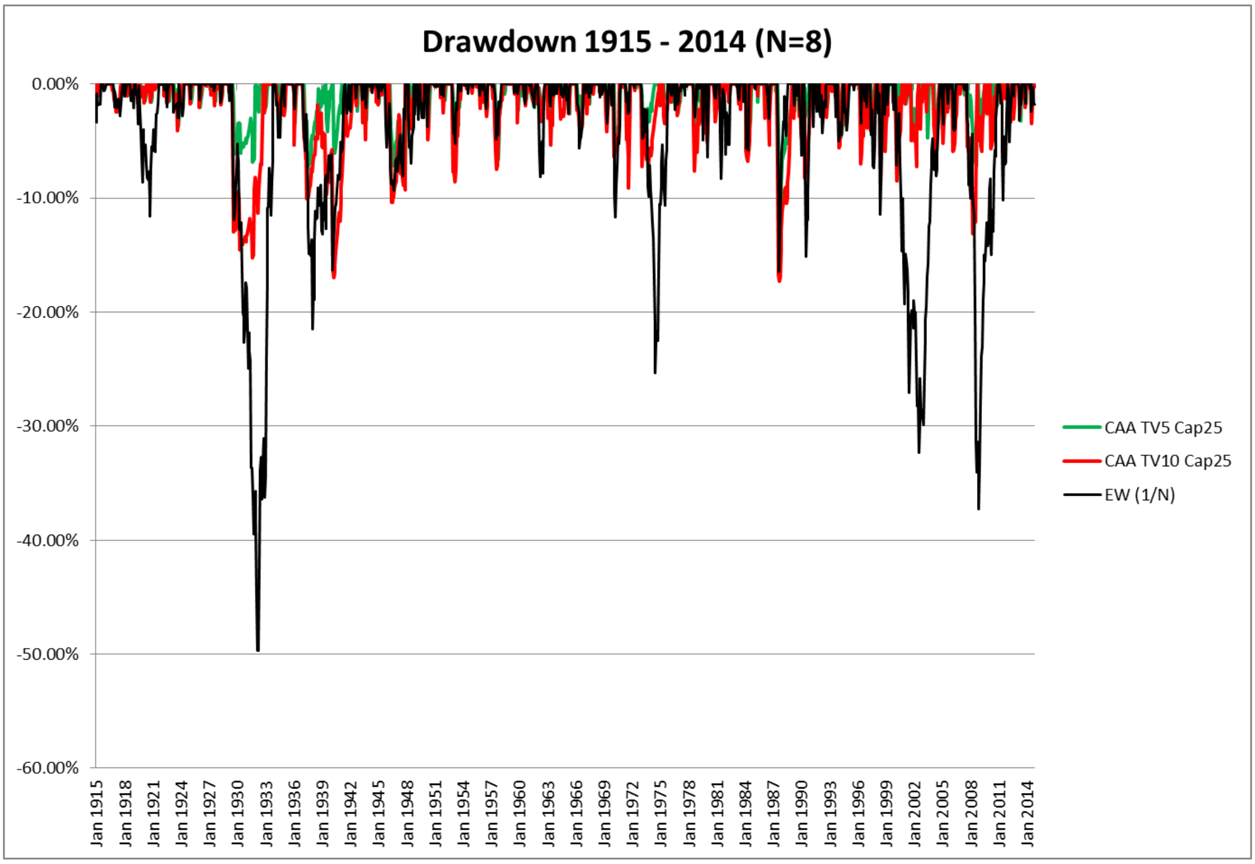

Fig. 9 Drawdowns for N=8

When we examine the portfolio weights in times of crisis, like the Great Depression or the Global Financial Crisis of 2008 (see Appendix A), you will see that CAA neatly switches to partly or full cash (here T-Bills and USGov10y) at intuitive times. This is just the result from the momentum effect (short term lookbacks) and the long-only MVO with no caps on cash.

8. The intermediate sized universe (N=16)

Here we will present the results for CAA for the intermediate size (N=16) universe for Jan 1914-Dec 2014 (100 years). This universe consists of the 10 Fama/French US sectors plus five US bonds: US Gov10y, US Gov30y, US Muni, US Corp, US High Yield and (3m) T-Bills. So it contains only US stocks and bonds. For some assets data are missing at the start and/or end of the century (see Keller 2014b). All non-cash assets will have the default 25% cap, while both cash assets (T-Bills and US10yGov) are uncapped. The benchmark is the equal weight (EW or 1/N) model of the full portfolio. The results for the offensive (TV=10%) and defensive (TV=5%) models are shown in fig. 10 and 12, resp. and in fig. 11 and 13.

TV=10% (N=16)	CAA	EW (1/N)
R=	11.20%	8.70%
V=	9.40%	11.50%
D=	19.70%	64.70%
SR5=	65.60%	32.70%
CR5=	31.30%	5.80%

Fig. 10 The offensive (TV=10%) results for N=16 (Jan 1915 – Dec 2014)

TV=5% (N=16)	CAA	EW (1/N)
R=	8.70%	8.70%
V=	5.90%	11.50%
D=	12.20%	64.70%
SR5=	62.50%	32.70%
CR5=	30.50%	5.80%

Fig. 11 The defensive (TV=5%) results for N=16 (Jan 1915 – Dec 2014)

As we see from the results for the N=16 universe, the returns R for CAA are less than in the smaller (N=8) universe. We suspect this is because the US-only universe is much less diversified relative to the global N=8 universe. The improvements for the N=16 CAA models relative to the benchmark are to be found in the return/risk statistics SR5 and CR5 which are again 100% and 500% better than EW for both the offensive as well as the defensive model. The improvement of the maximum drawdown D is again impressive, which is 1/5 and 1/3 of the EW value for the defensive and offensive model, resp.

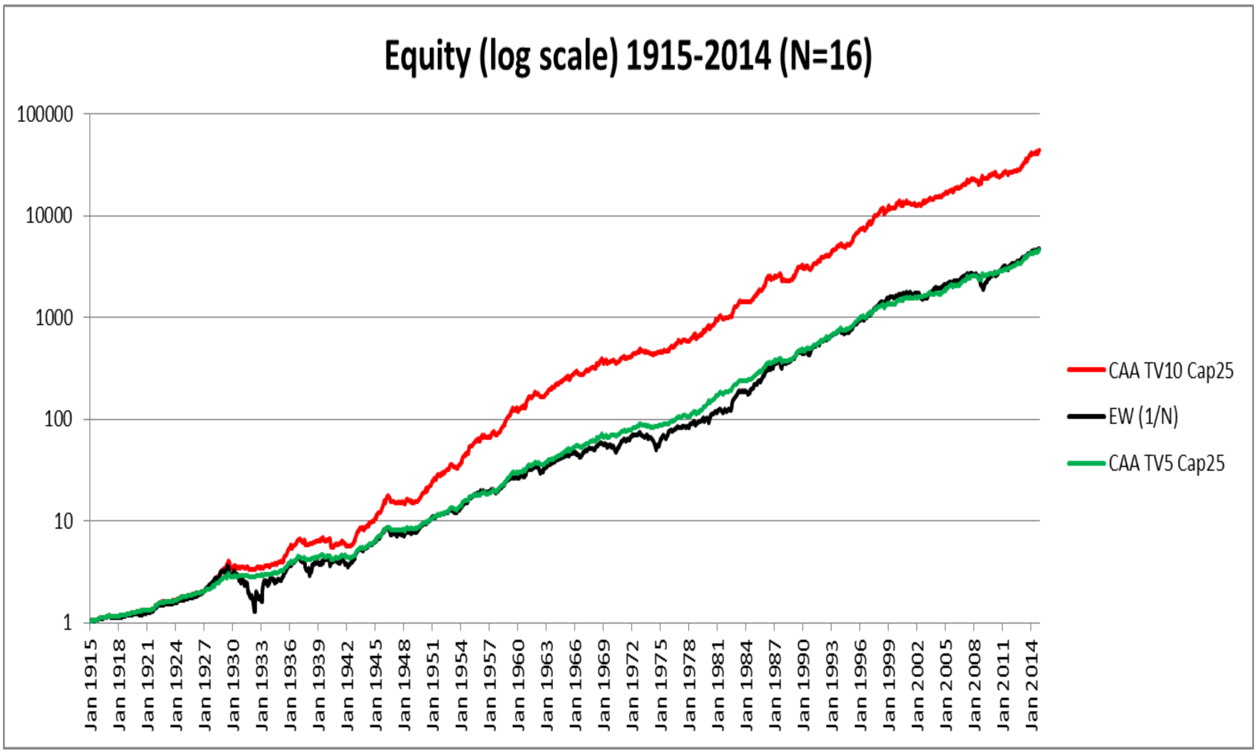

Fig. 12 Equity lines for N=16

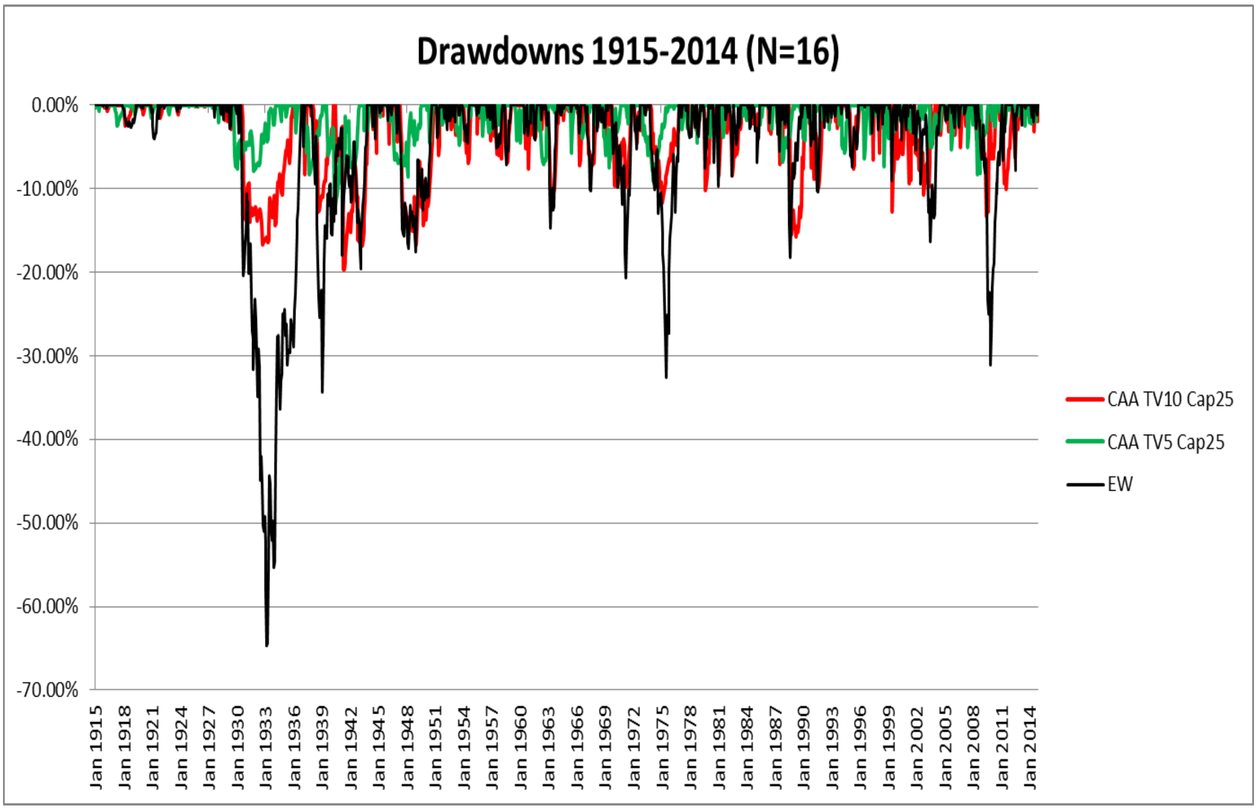

Fig. 13 Drawdowns for N=16

9. The large universe (N=39)

Here we will present the results for CAA for the large size (N=39) universe for Jan 1914-Dec 2014 (100 years). The universe consists of all the assets of the above two universes (N=8 and N=16) plus US Small Caps equities, GSCI, Gold, Foreign bonds, US TIPS, US Composite REITs, US Mortgage REITs, FTSE US 1000/US 1500/Global ex US/Developed/EM, JapanGov10y, Dow Util/Transport/Industry, FX-1x/2x, and Timber. For some assets data are missing at the start and end of the century (see Keller 2014b). All non-cash assets will have again the default 25% cap, while both cash assets (T-Bills and US10yGov) are uncapped. The benchmark is the equal weight (EW or 1/N) model of the full portfolio. The results for the offensive (TV=10%) and the defensive model (TV=5%) models are shown in fig. 14 and 16, resp. and fig. 17 and 18.

TV=10% (N=39)	CAA	EW (1/N)
R=	15.40%	8.80%
V=	10.40%	10.70%
D=	22.80%	63.30%
SR5=	100.20%	35.00%
CR5=	45.80%	5.90%

Fig. 14 The offensive (TV=10%) results for N=39 (Jan 1915 – Dec 2014)

TV=5% (N=39)	CAA	EW (1/N)
R=	11.80%	8.80%
V=	7.30%	10.70%
D=	15.60%	63.30%
SR5=	92.40%	35.00%
CR5=	43.30%	5.90%

Fig. 15 The defensive (TV=5%) results for N=39 (Jan 1915 – Dec 2014)

These results clearly demonstrate that MVO using CLA is robust to large universes where N>>T (such as N=16 and N=39 where T<=12). The results are impressive, in particular the return R for the offensive model, which is nearly double that of the EW benchmark. Its Sharpe ratio SR5 is above 1 (nearly three times the SR5 of the EW), while the Calmar ratio (CR5=46%) is more than seven times that of the EW benchmark. Its drawdown D=23% is substantial, although still nearly three times smaller than EW. Similar observations hold for the defensive model with lower V and D and nearly as good SR5 and CR5 as the offensive model has. Both models clearly beat EW in all respects.

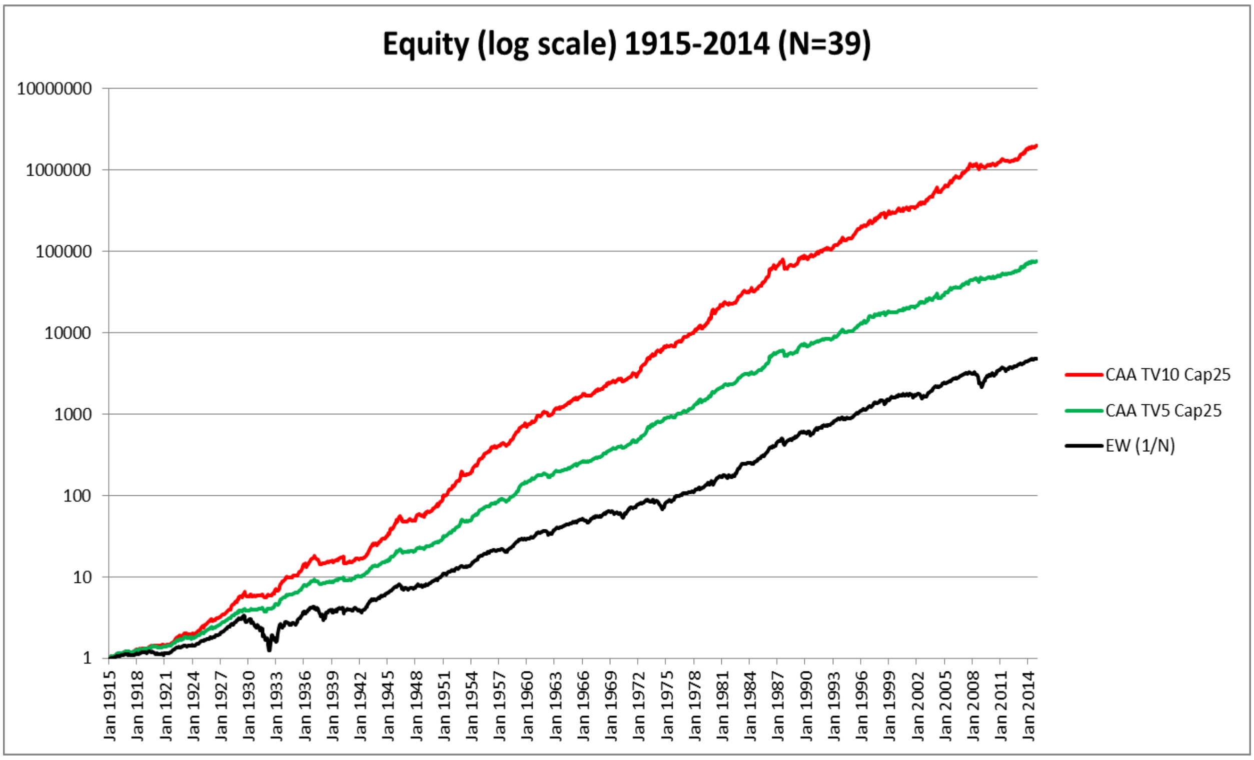

Fig. 16 Equity lines for N=39 (Jan 1915 – Dec 2014)

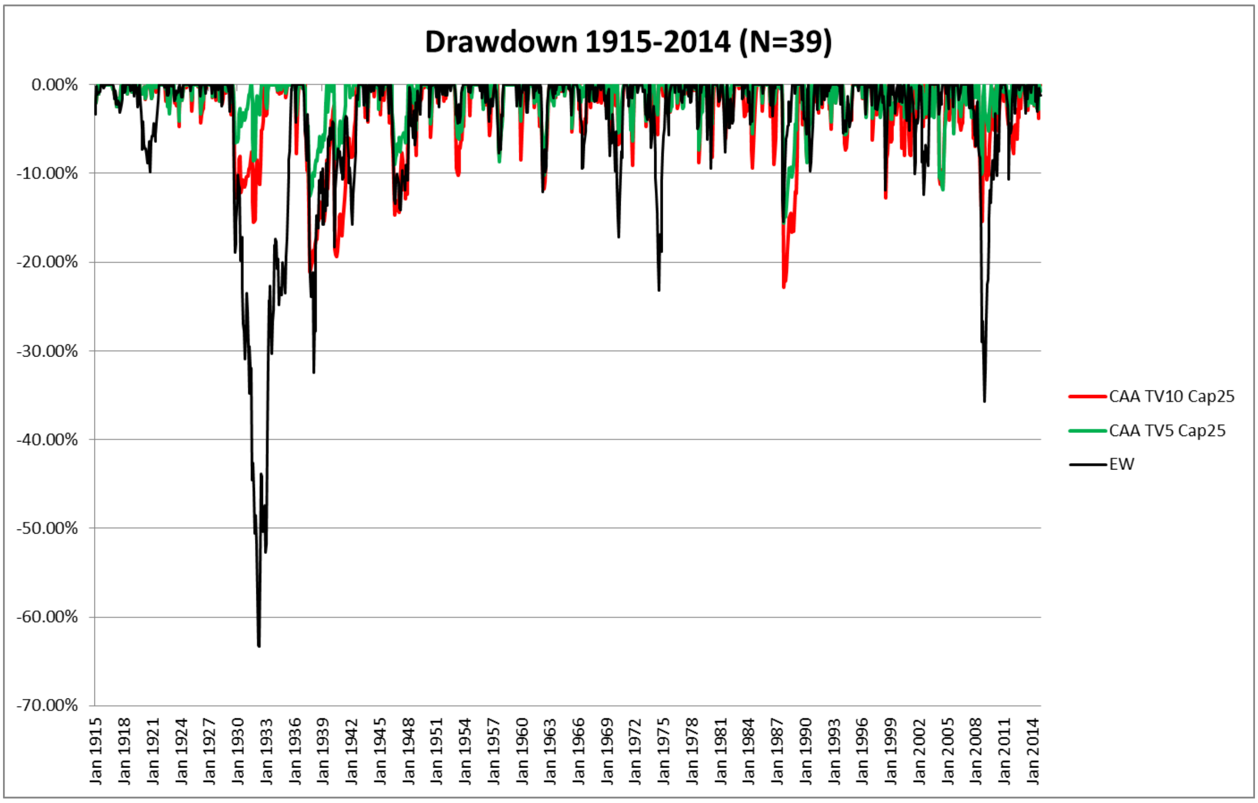

Fig. 17 Drawdowns for N=39 (Jan 1915 – Dec 2014)

10. Conclusions

When the first author submitted a paper (see Keller 2014a) to a great American Investment contest, he received the following reply from the chairman of the jury in response to his question about why he was not selected to place in the Top 3: “My guess is that some were turned off by the Markowitz affirmation”. So it seems Markowitz’ MPT is not in vogue any more in the investment industry. This is confirmed by the quotes (referred to above) made in the recent seminal book of Ang (2014).

In the present quant “blog scene”, Markowitz is also not very popular, as witnessed by this recent article from Wesley Gray (a very well-respected quant, also by us!):

> “Although Markowitz did win a Nobel Prize, and this was partly based on his elegant mathematical solution to identifying mean-variance efficient portfolios, a funny thing happened when his ideas were applied in the real world: **mean-variance performed poorly**. The fact that a Nobel-Prize winning idea translated into a no-value-add-situation for investors is something to keep in mind when considering any optimization method for asset allocation ...**complexity does not equal value!”**

This is the general theme in the investment industry: MVO is too complex, unstable, and unreliable. We disagree. It is the industry’s *implementation* of MVO (with long multi-year lookbacks and noisy short sales), not the MVO *theory*, which is incorrect.

In particular, using the last 60 month (5 year) mean return to forecast future returns runs directly counter to the theory of long-term price reversals. See also Asness (2012, 2015) who emphasized demonstrated that *negative* momentum over a 5 years lookback as a good proxy for value! And when our implementation works with even a 4 month estimate of the covariance matrix (as our Harry had shown) the culprit is probably not the covariance matrix but the combination of the wrong return prediction in combination with allowance for short sales. When we simply restrict our model to long-only and use momentum to estimate returns, we easily beat 1/N in all cases considered by a large margin.

One of the most charming characteristics (in our opinion) of combining MVO with momentum and long-only constraints is the nearly complete absence of degrees of freedom for datasnooping. We played with simple caps for reasons of practical implementation, but as we show in Appendix A, the dominance of our *Classical Asset Allocation* (CAA) model over 1/N is independent of the cap chosen, including no caps at all (unconstrained long-only). Our lookback horizon was deliberately chosen to span the known effective momentum space rather than optimizing on a specific lookback value. The same holds for the simple historical covariance matrix based on the past 12 months.

Our unpublished analyses (Butler, 2015) indicate that perturbing these parameters would cause no difference to our conclusions. So in terms of datasnooping, our results are robust. And because we tested over a century of data, these results also hold when rates are rising (as many expect after QE), as witnessed by our CAA results in the thirties (with zero yields) and during the seventies.

As suggestions for further research we plan to repeat our exercise with daily data and even allow for short sales and transactions costs[9]. Butler (2015, forthcoming) performs a Monte-Carlo style analysis,

[9] The yearly turnover of CAA for the various universes (with TV=10%/5%, cap=25%) were roughly between 4 (for N=8) and 7 (N=39). So, with (today's) transaction costs of less than 0.1% (and sometimes even zero), the annual return R (CAGR) for e.g. N=39 will in worst case be 0.7% lower after taking into account these transaction costs. Of course, in the early years of the century these transaction cost were much higher than

including random universes and random rebalance dates in order to observe a distribution of results for robust statistical tests. From this study we also conclude that as long as Markowitz and momentum are combined, MVO beats 1/N handily with very strong statistical significance, and also beats more robust return-agnostic optimizations like Equal Risk Contribution.

0.1%. With e.g. a difference of 14.4-8.8= 5.6% in R (for the N=39/TV=10/cap25% results) this implies that the average transaction costs over a century can be as high as 0.8% (= 5.6/7) before CAA loses on R from 1/N.

Appendix A. Some robustness tests

To make sure we did not pick the best value for the default cap (25%), we present in fig. 18 the results for the three universes for several other cap values, from 10%, 17%, 25% (default), 33%, 40%, 50% to 100% (=no caps), for both the offensive (TV=10%) and defensive (TV=5%) models. A cap=25% implies that there should be at least 4 assets in the portfolio assuming that there is no cash (uncapped) selected. This is relatively high for an N=8 universe (50%), so in practice one might choose a slightly higher cap, like cap33 or cap25.n For the N=39 universe the default cap of 25% seems a better fit. The choice of cap is a function of investor preferences for minimum thresholds of diversification and it is not presented for parameter optimization.

Below, we highlight in grey the best results in terms of caps. As can be seen from both tables, our default cap of 25% is clearly not always optimal in terms of return. Notice that we get the best performance (R=16.3%) with N=39, TV=10% and caps=50% with slightly better SR5 and CR5 than with the default cap=25% (R=15.4%). The worst performance (R=8.1%) for N=8, TV=5 and cap=10%, is slightly less than EW (8.7%) but still with much better return/risk (SR5 nearly doubled and CR5 is five times EW) and with the best drawdown (D=10.7%) overall. The worst drawdown overall (D=23.2%) for N=16, TV5, cap10% is still nearly three times better than EW (64.7%), while the Sharpe ratio (SR5=60.7%) is still nearly double that of EW (32.7%). The worst Calmar ratio (CR5=27.7% for N=16, TV10, cap 10%) is still nearly five times better than EW. So CAA is much better than EW, independent of the choice of cap (and the choice of universe and TV).

Notice also that cap level impacts realized ex post volatility V relative to target volatility (TV). In general, (but with small exceptions), higher caps result in higher ex post volatility V. For N=8 and TV=10% the best corresponding (to TV) is cap 50 % (with a V=10.1%); for N=16 it is cap 40% (V=10.1%); for N=39 it is cap 17% (V=9.8%). These results are intuitive, as smaller caps on a larger universe will still allow for highly concentrated portfolios relative to the size of the overall universe.

In fig. 19 we show the optimal weights for N=8, TV=10%, Cap=25% for the years 2007-2009. Notice how the optimal allocation switches to cash as a result of momentum and changes in covariances long before the Lehman crisis (in Sep 2008). Also note how the cap of 25% leads to nearly EW like (but not 1/N) optimal portfolios in 2007, which reduces turnover.

Lastly (in fig. 20) we show four examples of EFs, where the corresponding “degeneration” of the number of corner points (from 7 in 2007 to only two in Sep 2008) becomes clear. The explanation is simple: with so many assets with negative returns before and during the crisis in 2008, only a limited number of assets with positive returns (mostly bonds, but also some Tech and EEM before Jun 2008, see fig. 19) can be included, which limits the number of corner points.

N8 TV10	Cap10	Cap17	Cap25	Cap33	Cap40	Cap50	Cap100%	EW (1/N)
R=	8.80%	11.00%	12.70%	13.60%	14.10%	14.60%	15.20%	8.70%
V=	5.30%	6.90%	8.30%	9.20%	9.60%	10.10%	11.30%	9.20%
D=	-11.30%	-15.80%	-17.30%	-20.50%	-21.80%	-21.80%	-21.80%	-49.70%
SR5=	71.80%	86.20%	92.30%	93.50%	94.20%	94.90%	90.20%	40.10%
CR5=	33.50%	37.70%	44.60%	42.00%	41.60%	44.10%	46.80%	7.40%
N8 TV5	**Cap10**	**Cap17**	**Cap25**	**Cap33**	**Cap40**	**Cap50**	**Cap100%**	**EW (1/N)**
R=	8.30%	9.60%	10.50%	10.80%	10.90%	11.10%	11.00%	8.70%
V=	4.40%	5.20%	5.80%	6.10%	6.30%	6.50%	6.70%	9.20%
D=	-8.10%	-10.10%	-10.60%	-10.60%	-10.70%	-13.00%	-17.20%	-49.70%
SR5=	76.00%	88.70%	94.50%	94.80%	94.40%	95.10%	89.60%	40.10%
CR5=	40.70%	45.80%	51.50%	54.70%	55.70%	47.20%	35.10%	7.40%

N16 TV10	Cap10	Cap17	Cap25	Cap33	Cap40	Cap50	Cap100%	EW (1/N)
R=	9.80%	10.60%	11.20%	11.50%	11.70%	12.00%	12.10%	8.70%
V=	7.70%	8.70%	9.40%	9.80%	10.10%	10.40%	10.80%	11.50%
D=	-17.20%	-18.90%	-19.70%	-19.30%	-19.50%	-21.80%	-23.20%	-64.70%
SR5=	61.80%	64.20%	65.60%	66.40%	67.20%	67.10%	65.20%	32.70%
CR5=	27.70%	29.60%	31.30%	33.70%	34.60%	31.90%	30.50%	5.80%
N16 TV5	**Cap10**	**Cap17**	**Cap25**	**Cap33**	**Cap40**	**Cap50**	**Cap100%**	**EW (1/N)**
R=	8.10%	8.60%	8.70%	8.70%	8.80%	8.80%	8.80%	8.70%
V=	5.20%	5.60%	5.90%	6.10%	6.20%	6.20%	6.30%	11.50%
D=	10.70%	11.90%	12.20%	11.30%	11.40%	12.00%	12.60%	64.70%
SR5=	60.70%	63.90%	62.50%	61.60%	60.90%	61.30%	60.10%	32.70%
CR5=	29.30%	30.40%	30.50%	33.00%	32.80%	31.60%	29.90%	5.80%

N39 TV10	Cap10	Cap17	Cap25	Cap33	Cap40	Cap50	Cap100%	EW (1/N)
R=	13.00%	14.70%	15.40%	16.00%	16.20%	16.30%	15.70%	8.80%
V=	8.80%	9.80%	10.40%	11.00%	11.40%	11.90%	12.90%	10.70%
D=	-18.70%	-21.40%	-22.80%	-23.60%	-23.50%	-23.10%	-22.70%	-63.30%
SR5=	90.40%	98.90%	100.20%	100.10%	98.20%	95.50%	83.60%	35.00%
CR5=	42.70%	45.00%	45.80%	46.70%	47.70%	49.10%	47.40%	5.90%
N39 TV5	**Cap10**	**Cap17**	**Cap25**	**Cap33**	**Cap40**	**Cap50**	**Cap100%**	**EW (1/N)**
R=	10.50%	11.40%	11.80%	12.00%	12.10%	12.10%	11.90%	8.80%
V=	6.20%	6.90%	7.30%	7.60%	7.80%	8.00%	8.10%	10.70%
D=	-11.70%	-15.90%	-15.60%	-15.30%	-15.30%	-15.30%	-16.10%	-63.30%
SR5=	88.70%	93.30%	92.40%	92.60%	91.20%	89.00%	85.50%	35.00%
CR5=	47.40%	40.50%	43.30%	46.10%	46.40%	46.40%	43.30%	5.90%

Fig. 18 Results various caps for N=8, 16, 39 and TV=10, 5% (Jan 1915 – Dec 2014)

N8 T10	T-Bill	US10yr	HYield	SP500	Tech	EAFE	JapTopix	EEM
2007-01-31	0%	0%	3%	25%	22%	25%	0%	25%
2007-02-28	0%	0%	25%	0%	0%	25%	25%	25%
2007-03-31	0%	0%	25%	25%	0%	25%	0%	25%
2007-04-30	0%	0%	0%	25%	25%	25%	0%	25%
2007-05-31	0%	0%	0%	25%	25%	25%	0%	25%
2007-06-30	0%	0%	0%	25%	25%	25%	0%	25%
2007-07-31	0%	0%	0%	25%	25%	25%	0%	25%
2007-08-31	0%	0%	0%	25%	25%	25%	0%	25%
2007-09-30	0%	0%	0%	25%	25%	25%	0%	25%
2007-10-31	0%	0%	0%	25%	25%	25%	0%	25%
2007-11-30	0%	75%	0%	0%	0%	0%	0%	25%
2007-12-31	0%	75%	0%	0%	0%	0%	0%	25%
2008-01-31	0%	100%	0%	0%	0%	0%	0%	0%
2008-02-29	0%	75%	0%	0%	0%	0%	0%	25%
2008-03-31	0%	100%	0%	0%	0%	0%	0%	0%
2008-04-30	0%	75%	0%	0%	0%	0%	0%	25%
2008-05-31	0%	25%	25%	0%	25%	0%	0%	25%
2008-06-30	0%	100%	0%	0%	0%	0%	0%	0%
2008-07-31	0%	100%	0%	0%	0%	0%	0%	0%
2008-08-31	0%	100%	0%	0%	0%	0%	0%	0%
2008-09-30	0%	100%	0%	0%	0%	0%	0%	0%
2008-10-31	0%	100%	0%	0%	0%	0%	0%	0%
2008-11-30	8%	92%	0%	0%	0%	0%	0%	0%
2008-12-31	12%	88%	0%	0%	0%	0%	0%	0%
2009-01-31	22%	78%	0%	0%	0%	0%	0%	0%
2009-02-28	21%	79%	0%	0%	0%	0%	0%	0%
2009-03-31	22%	78%	0%	0%	0%	0%	0%	0%
2009-04-30	0%	75%	25%	0%	0%	0%	0%	0%
2009-05-31	51%	12%	25%	0%	0%	0%	0%	12%
2009-06-30	59%	0%	25%	0%	16%	0%	0%	0%
2009-07-31	46%	18%	25%	0%	0%	0%	0%	11%
2009-08-31	20%	39%	25%	0%	15%	0%	0%	0%
2009-09-30	34%	31%	25%	0%	0%	0%	0%	11%
2009-10-31	30%	32%	25%	0%	0%	0%	0%	13%
2009-11-30	0%	40%	25%	0%	20%	0%	0%	15%
2009-12-31	32%	0%	25%	0%	25%	0%	0%	18%

Fig. 19 Weights for N=8 and TV=10% (Jan 2007 – Dec 2009)

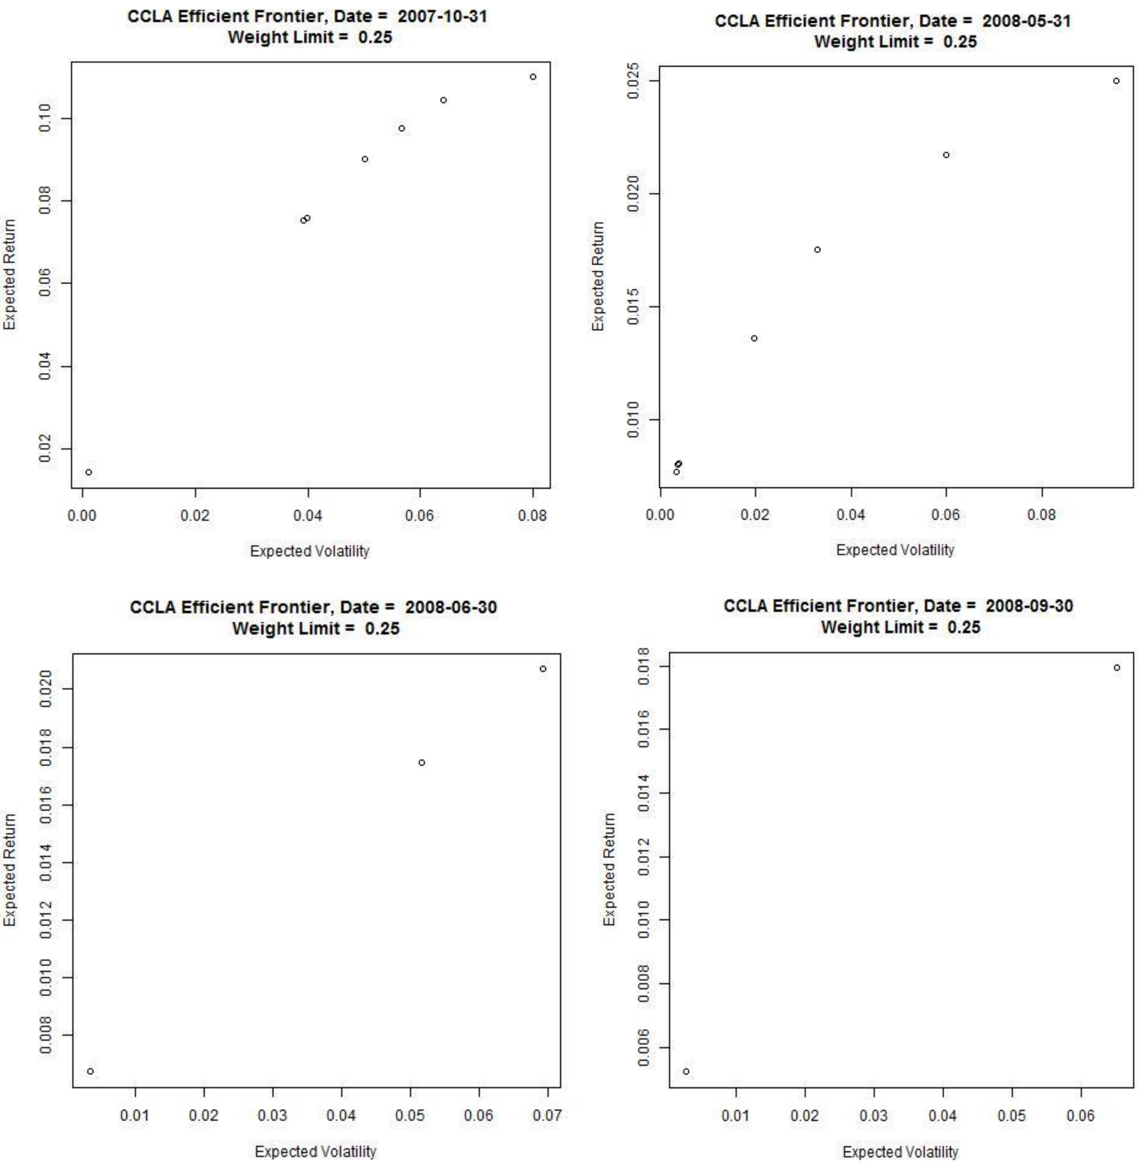

Fig. 20 Four Efficient Frontiers for N=8 and TV=10% (Oct 2007, and May, Jun, Sep 2008)

Appendix B. CLA code (in R) by Ilya Kipnis (QuantStratTradeR[10])

```
require(quantmod)
require(PerformanceAnalytics)
require(TTR)

CCLA <- function(covMat, retForecast, maxIter = 1000,
            verbose = FALSE, scale = 252,
            weightLimit = .7, volThresh = .1) {
 if(length(retForecast) > length(unique(retForecast))) {
  sequentialNoise <- seq(1:length(retForecast)) * 1e-12
  retForecast <- retForecast + sequentialNoise
 }

 #initialize original out/in/up status
 if(length(weightLimit) == 1) {
  weightLimit <- rep(weightLimit, ncol(covMat))
 }
 rankForecast <- length(retForecast) - rank(retForecast) + 1
 remainingWeight <- 1 #have 100% of weight to allocate
 upStatus <- inStatus <- rep(0, ncol(covMat))
 i <- 1
 while(remainingWeight > 0) {
  securityLimit <- weightLimit[rankForecast == i]
  if(securityLimit < remainingWeight) {
   upStatus[rankForecast == i] <- 1 #if we can't invest all remaining weight into the security
   remainingWeight <- remainingWeight - securityLimit
  } else {
   inStatus[rankForecast == i] <- 1
   remainingWeight <- 0
  }
  i <- i + 1
 }

 #initial matrices (W, H, K, identity, negative identity)
 covMat <- as.matrix(covMat)
 retForecast <- as.numeric(retForecast)
 init_W <- cbind(2*covMat, rep(-1, ncol(covMat)))
 init_W <- rbind(init_W, c(rep(1, ncol(covMat)), 0))
 H_vec <- c(rep(0, ncol(covMat)), 1)
 K_vec <- c(retForecast, 0)
 negIdentity <- -1*diag(ncol(init_W))
 identity <- diag(ncol(init_W))
 matrixDim <- nrow(init_W)
 weightLimMat <- matrix(rep(weightLimit, matrixDim), ncol=ncol(covMat), byrow=TRUE)

 #out status is simply what isn't in or up
 outStatus <- 1 - inStatus - upStatus
```

[10] https://quantstrattrader.wordpress.com/

```
  #initialize expected volatility/count/turning points data structure
  expVol <- Inf
  lambda <- 100
  count <- 0
  turningPoints <- list()
  while(lambda > 0 & count < maxIter) {

    #old lambda and old expected volatility for use with numerical algorithms
    oldLambda <- lambda
    oldVol <- expVol

    count <- count + 1

    #compute W, A, B
    inMat <- matrix(rep(c(inStatus, 1), matrixDim), nrow = matrixDim, byrow = TRUE)
    upMat <- matrix(rep(c(upStatus, 0), matrixDim), nrow = matrixDim, byrow = TRUE)
    outMat <- matrix(rep(c(outStatus, 0), matrixDim), nrow = matrixDim, byrow = TRUE)

    W <- inMat * init_W + upMat * identity + outMat * negIdentity

    inv_W <- solve(W)
    modified_H <- H_vec - rowSums(weightLimMat* upMat[,-matrixDim] * init_W[,-matrixDim])
    A_vec <- inv_W %*% modified_H
    B_vec <- inv_W %*% K_vec

    #remove the last elements from A and B vectors
    truncA <- A_vec[-length(A_vec)]
    truncB <- B_vec[-length(B_vec)]

    #compute in Ratio (aka Ratio(1) in Kwan.xls)
    inRatio <- rep(0, ncol(covMat))
    inRatio[truncB > 0] <- -truncA[truncB > 0]/truncB[truncB > 0]

    #compute up Ratio (aka Ratio(2) in Kwan.xls)
    upRatio <- rep(0, ncol(covMat))
    upRatioIndices <- which(inStatus==TRUE & truncB < 0)
    if(length(upRatioIndices) > 0) {
      upRatio[upRatioIndices] <- (weightLimit[upRatioIndices] - truncA[upRatioIndices]) /
truncB[upRatioIndices]
    }

    #find lambda -- max of up and in ratios
    maxInRatio <- max(inRatio)
    maxUpRatio <- max(upRatio)
    lambda <- max(maxInRatio, maxUpRatio)

    #compute new weights
    wts <- inStatus*(truncA + truncB * lambda) + upStatus * weightLimit + outStatus * 0

    #compute expected return and new expected volatility
    expRet <- t(retForecast) %*% wts
```

```
  expVol <- sqrt(wts %*% covMat %*% wts) * sqrt(scale)

  #create turning point data row and append it to turning points
  turningPoint <- cbind(count, expRet, lambda, expVol, t(wts))
  colnames(turningPoint) <- c("CP", "Exp. Ret.", "Lambda", "Exp. Vol.", colnames(covMat))
  turningPoints[[count]] <- turningPoint

  #binary search for volatility threshold -- if the first iteration is lower than the threshold,
  #then immediately return, otherwise perform the binary search until convergence of lambda
  if(oldVol == Inf & expVol < volThresh) {
   turningPoints <- do.call(rbind, turningPoints)
   threshWts <- tail(turningPoints, 1)
   return(list(turningPoints, threshWts))
  } else if(oldVol > volThresh & expVol < volThresh) {
   upLambda <- oldLambda
   dnLambda <- lambda
   meanLambda <- (upLambda + dnLambda)/2
   while(upLambda - dnLambda > .00001) {

    #compute mean lambda and recompute weights, expected return, and expected vol
    meanLambda <- (upLambda + dnLambda)/2
    wts <- inStatus*(truncA + truncB * meanLambda) + upStatus * weightLimit + outStatus * 0
    expRet <- t(retForecast) %*% wts
    expVol <- sqrt(wts %*% covMat %*% wts) * sqrt(scale)

    #if new expected vol is less than threshold, mean becomes lower bound
    #otherwise, it becomes the upper bound, and loop repeats
    if(expVol < volThresh) {
     dnLambda <- meanLambda
    } else {
     upLambda <- meanLambda
    }
   }

   #once the binary search completes, return those weights, and the corner points
   #computed until the binary search. The corner points aren't used anywhere, but they're there.
   threshWts <- cbind(count, expRet, meanLambda, expVol, t(wts))
   colnames(turningPoint) <- colnames(threshWts) <- c("CP", "Exp. Ret.", "Lambda", "Exp. Vol.",
colnames(covMat))
   turningPoints[[count]] <- turningPoint
   turningPoints <- do.call(rbind, turningPoints)
   return(list(turningPoints, threshWts))
  }

  #this is only run for the corner points during which binary search doesn't take place
  #change status of security that has new lambda
  if(maxInRatio > maxUpRatio) {
   inStatus[inRatio == maxInRatio] <- 1 - inStatus[inRatio == maxInRatio]
   upStatus[inRatio == maxInRatio] <- 0
  } else {
   upStatus[upRatio == maxUpRatio] <- 1 - upStatus[upRatio == maxUpRatio]
   inStatus[upRatio == maxUpRatio] <- 0
```

```
    }
    outStatus <- 1 - inStatus - upStatus
  }

  #we only get here if the volatility threshold isn't reached
  #can actually happen if set sufficiently low
  turningPoints <- do.call(rbind, turningPoints)

  threshWts <- tail(turningPoints, 1)

  return(list(turningPoints, threshWts))
}

sumIsNa <- function(column) {
  return(sum(is.na(column)))
}

returnForecast <- function(prices) {
  forecast <- (ROC(prices, n = 1, type="discrete") + ROC(prices, n = 3, type="discrete") +
               ROC(prices, n = 6, type="discrete") + ROC(prices, n = 12, type="discrete"))/22
  forecast <- as.numeric(tail(forecast, 1))
  return(forecast)
}

kellerCLAfun <- function(prices, returnWeights = FALSE,
                         weightLimit, volThresh, uncappedAssets) {

  if(sum(colnames(prices) %in% uncappedAssets) == 0) {
    stop("No assets are uncapped.")
  }
  #initialize data structure to contain our weights
  weights <- list()

  #compute returns
  returns <- Return.calculate(prices)
  returns[1,] <- 0 #impute first month with zeroes
  ep <- endpoints(returns, on = "months")
  for(i in 2:(length(ep) - 12)) {
    priceSubset <- prices[ep[i]:ep[i+12]] #subset prices
    retSubset <- returns[ep[i]:ep[i+12]] #subset returns
    assetNAs <- apply(retSubset, 2, sumIsNa)
    zeroNAs <- which(assetNAs == 0)
    priceSubset <- priceSubset[, zeroNAs]
    retSubset <- retSubset[, zeroNAs]

    #remove perfectly correlated assets
    retCors <- cor(retSubset)
    diag(retCors) <- NA
    corMax <- round(apply(retCors, 2, max, na.rm = TRUE), 7)
    while(max(corMax) == 1) {
      ones <- which(corMax == 1)
      valid <- which(!names(corMax) %in% uncappedAssets)
```

```
      toRemove <- intersect(ones, valid)
      toRemove <- max(valid)
      retSubset <- retSubset[, -toRemove]
      priceSubset <- priceSubset[, -toRemove]
      retCors <- cor(retSubset)
      diag(retCors) <- NA
      corMax <- round(apply(retCors, 2, max, na.rm = TRUE), 7)
    }

    covMat <- cov(retSubset) #compute covariance matrix

    #Dr. Keller's return forecast
    retForecast <- returnForecast(priceSubset)
    uncappedIndex <- which(colnames(covMat) %in% uncappedAssets)
    weightLims <- rep(weightLimit, ncol(covMat))
    weightLims[uncappedIndex] <- 1

    cla <- CCLA(covMat = covMat, retForecast = retForecast, scale = 12,
                weightLimit = weightLims, volThresh = volThresh) #run CCLA algorithm
    CPs <- cla[[1]] #corner points
    wts <- cla[[2]] #binary search volatility targeting -- change this line and the next
    #if using max sharpe ratio golden search
    wts <- wts[, 5:ncol(wts)] #from 5th column to the end
    if(length(wts) == 1) {
      names(wts) <- colnames(covMat)
    }

    zeroes <- rep(0, ncol(prices) - length(wts))
    names(zeroes) <- colnames(prices)[!colnames(prices) %in% names(wts)]
    wts <- c(wts, zeroes)
    wts <- wts[colnames(prices)]

    #append to weights
    wts <- xts(t(wts), order.by=tail(index(retSubset), 1))
    weights[[i]] <- wts
  }
  weights <- do.call(rbind, weights)

  #compute strategy returns
  stratRets <- Return.portfolio(returns, weights = weights)
  if(returnWeights) {
    return(list(weights, stratRets))
  }
  return(stratRets)
}

kellerCLAcomparison <- function(prices, weightLimits, volThresh,
                                uncappedAssets) {
  #initialize configurations
  configs <- list()
  for(i in 1:length(weightLimits)) {
```

```
    weightLimit <- weightLimits[i]

    #run CLA function loop -- because of differing numbers of securities,
    #it'll take care of weight limits inside new function
    config <- kellerCLAfun(prices = prices, returnWeights = FALSE,
                 weightLimit = weightLimit, volThresh = volThresh,
                 uncappedAssets = uncappedAssets)

    #create appropriate column name for configuration
    colnames(config) <- paste("thresh", volThresh*100, "wLim", round(weightLimits[i], 2)*100,
sep="_")

    #append to configurations
    configs[[i]] <- config
  }

  #combine configurations
  configs <- do.call(cbind, configs)
  return(configs)
}

#plot efficient frontier
claEfPlot <- function(prices, weightLimit = .5, uncappedAssets = uncappedAssets) {
  returns <- Return.calculate(prices)
  returns[1,] <- 0
  ep <- endpoints(prices, on = "months")

  for(i in 2:(length(ep)-12)) {
    priceSubset <- prices[ep[i]:ep[i+12]]
    retSubset <- returns[ep[i]:ep[i+12]]
    assetNAs <- apply(retSubset, 2, sumIsNa)
    zeroNAs <- which(assetNAs == 0)
    priceSubset <- priceSubset[, zeroNAs]
    retSubset <- retSubset[, zeroNAs]

    retCors <- cor(retSubset)
    diag(retCors) <- NA
    corMax <- round(apply(retCors, 2, max, na.rm = TRUE), 7)
    while(max(corMax) == 1) {
      ones <- which(corMax == 1)
      valid <- which(!names(corMax) %in% uncappedAssets)
      toRemove <- intersect(ones, valid)
      toRemove <- max(valid)
      retSubset <- retSubset[, -toRemove]
      priceSubset <- priceSubset[, -toRemove]
      retCors <- cor(retSubset)
      diag(retCors) <- NA
      corMax <- round(apply(retCors, 2, max, na.rm = TRUE), 7)
    }

    retForecast <- returnForecast(priceSubset)
    covMat <- cov(retSubset)
```

```
    uncappedIndex <- which(colnames(covMat) %in% uncappedAssets)
    weightLims <- rep(weightLimit, ncol(covMat))
    weightLims[uncappedIndex] <- 1


    cla <- CCLA(covMat = covMat, retForecast = retForecast, scale = 12,
                weightLimit = weightLims, volThresh = 0) #run CCLA algorithm
    if(!"ef" %in% dir()) {
      dir.create("ef")
    }
    jpeg(filename = paste0("ef/",as.character(tail(index(priceSubset), 1)),"_", "Frontier.jpeg"))
    plot(cla[[1]][, 2]*sqrt(12)~cla[[1]][,4], xlab="Expected Volatility", ylab="Expected Return",
         main = paste("CCLA Efficient Frontier, Date = ", as.character(tail(index(priceSubset), 1)),
                      "\nWeight Limit = ", weightLimit))
    dev.off()
  }
}
```

Literature

Ang, A, 2014, Asset Management: A Systematic Approach to Factor Investing (Financial Management Association Survey and Synthesis Series), Oxford University Press. See also Ang, A., 2012, Mean-Variance Investing, SSRN 2131932

Antonacci, G, 2011, Optimal Momentum: A Global Cross Asset Approach, SSRN 1833722

Antonacci, G., 2013a, Absolute Momentum: A Simple Rule-Based Strategy and Universal Trend Following Overlay, SSRN 2244633

Antonacci, G., 2013b, Risk Premia Harvesting Through Dual Momentum SSRN 2042750

Asness, C.S, Moskowitz, T.J, and Pedersen, L.H., 2012, Value and Momentum Everywhere, Working Paper nr. 80, The Initiative on Global Markets, University of Chicago.

Asness, C.S., Frazzini, A., Israel, R. and Moskowitz, T.J., 2014, Fact, Fiction and Momentum Investing. Journal of Portfolio Management, Fall 2014 (40th Anniversary Issue); Fama-Miller Working Paper. SSRN 2435323

Asness, C. S., Frazzini, A., Israel, R. and Moskowitz, T. J., 2015, Fact, Fiction, and Value Investing SSRN 2595747

Bailey, D.H., and López de Prado, M., 2013, An Open-Source Implementation of the Critical-Line Algorithm for Portfolio Optimization, Algorithms 2013, 6, 169-196 (also SSRN 2197616)

Butler, A., Philbrick, M., and Gordillo, R., 2012, Adaptive Asset Allocation: A Primer, SSRN 2328254

Butler, A., Butler A., Keller, W.J., Philbrick M., and Gordillo M., 2015, Mean Variance Optimization for Adaptive Asset Allocation: A Statistical Validation (forthcoming)

Choueifaty, Y., Froidure, T., and Reynier, J., 2011, Properties of the Most Diversified Portfolio, SSRN 1895459

DeMiguel, V., Garlappi, L., and Uppal, R., 2009, Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? The Review of Financial Studies 22/5, pp.1915-1953

Faber, M. T., 2007, A Quantitative Approach to Tactical Asset Allocation, Journal of Wealth Management, Spring 2007. Update (2013) as SSRN 962461.

Faber, M. T., 2010, Relative Strength Strategies for Investing, 2010, SSRN: 1585517

Gray, W., 2014, http://blog.alphaarchitect.com/2014/12/02/our-robust-asset-allocation-raa-solution/

Hallerbach, W.G., 2013, Advances in Portfolio Risk Control. Risk ! Parity ?, SSRN 2259041

Hurst, B., Ooi, Y.H., and Pedersen, L.H., 2012, A Century of Evidence on Trend-Following Investing, working paper, AQR Capital Management. 21

Jacobs H., Müller, S., and Weber, M. , 2013, How should individual investors diversify? An empirical evaluation of alternative asset allocation policies, SSRN 1471955

Jegadeesh, N., and Titman, S., 1993, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, Journal of Finance XLVIII, 65/91.

Jurczenko, E., Michel, T., and Teiletche, J., 2013, Generalized Risk-Based Investing, SSRN 2205979

Keller, W.J. and Van Putten, H., 2012, Generalized Momentum and Flexible Asset Allocation (FAA): An Heuristic Approach, SSRN 2193735

Keller, W.J. and Van Putten, H., 2013, Tactical MPT and Momentum: The Modern Asset Allocation (MAA), SSRN 2373086

Keller, W.J., 2014a, Momentum, Markowitz, and Smart Beta, A Tactical, Analytical and Practical Look at Modern Portfolio Theory, SSRN 2450017

Keller, W.J. and Butler, A., 2014b, A Century of Generalized Momentum; From Flexible Asset Allocations (FAA) to Elastic Asset Allocation (EAA). Available at SSRN 2543979

Kritzman, M., Page, S., and Turkington, D., 2010, In Defense of Optimization: The Fallacy of 1/N, Financial Analysts Journal 66-2, pp. 31-39

Kwan, C. C. , 2007, A Simple Spreadsheet-Based Exposition of the Markowitz Critical Line Method for Portfolio Selection, Spreadsheets in Education (eJSiE): Vol. 2: Iss. 3, Article 2.

Ma, T. and Jagannathan, R., 2002, Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps (May 2002). NBER Working Paper No. w8922. SSRN 310469

Maillard, S., Roncalli, T., and Teiletche, J., 2008, On the properties of equally-weighted risk contributions portfolios, SSRN 1271972.

Markowitz, H. M. , 1952, Portfolio Selection, Journal of Finance, 7(1), 77–91.

Markowitz, H. M. , 1956, The Optimization of a Quadratic Function Subject to Linear Constraints, Naval Research Logistics Quarterly, III, 111–133.

Markowitz, H. M. , 1959, Portfolio Selection: Efficient Diversification of Investments. John Wiley and Sons, New York, and 1991 2nd ed., Basil Blackwell, Cambridge, MA.

Markowitz, H. M., and P. Todd (2000), Mean-Variance Analysis in Portfolio Choice and Capital Markets. Frank J. Fabozzi Associates, New Hope, Pennsylvania.

Michaud, R. (1998), Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization. Oxford University Press.

Moskowitz, T., Ooi, Y.H., and Pedersen, L.H., 2011, Time Series Momentum, Working Paper nr. 79, The Initiative on Global Markets, University of Chicago.

Nawrocki, D., 1996, Portfolio analysis with a large universe of assets, Applied Economics 28, pp. 1191-1198.

Niedermayer, A., Niedermayer, D., 2007, Applying Markowitz's Critical Line Algorithm, Discussion paper 06-02, University of Bern, Dept. of Economics (also 2006 version SSRN 894842).

Roncalli, T., 2013, Introduction to Risk Parity and Budgeting, Chapman & Hall/CRC Financial Mathematics Series.

Sharpe, W.F., 1963, A Simplified Model for Portfolio Analysis, Management Science, Vol. 9, No. 2, pp. 277-293