

Optimizing Omega*

Helmut Mausser[†] David Saunders[‡] Luis Seco[§]

1 Introduction

Performance measurement that accurately reflects the goals of fund investors and managers has long been a topic of active discussion. Most modern performance measures differ from classical ones (such as the Sharpe ratio) in two key ways; first, they reflect the market practice of assessing performance against a benchmark, second, they account for the asymmetry in returns distributions by separately considering upside and downside. The Omega is a recent measure introduced by Shadwick and Keating [2002] possessing both of these characteristics. It is defined, for an asset with return R (a random variable) and benchmark return L as:

$$\Omega(L) = \frac{\int_L^\infty (1 - F(x)) dx}{\int_{-\infty}^L F(x) dx} = \frac{E[\max(R - L, 0)]}{E[\max(L - R, 0)]} \quad (1)$$

where F is the cumulative probability distribution function of the asset's return, $F(x) = P[R \leq x]$. Other measures sharing with Omega the properties of assessing performance against a benchmark and accounting for asymmetry include the closely related regret-reward measures $D_\lambda(L) = E[\max(R - L, 0)] - \lambda \cdot E[\max(L - R, 0)]$ studied by Dembo and his colleagues (e.g. Dembo and Rosen [1999] and Dembo and Mausser [2000]) and the kappa ratios $\kappa_n(L) = E[R - L] / \sqrt[n]{E[\max(L - R, 0)^n]}$ introduced by Kaplan and Knowles [2003] as generalizations of the Sortino ratio (see, e.g. Sortino, Van der Meer and Plantinga [1999]).

*The authors are grateful to F. Glover, G. Comezaña and two anonymous referees for helpful comments and suggestions.

[†]Helmut Mausser is a senior mathematician at Algorithmics Inc. He may be reached at: hmausser@algorithmics.com

[‡]Corresponding Author: David Saunders is an assistant professor in the Department of Statistics and Actuarial Science at the University of Waterloo. He may be reached at: dsaunders@math.uwaterloo.ca

[§]Luis Seco is professor in the Department of Mathematics and the Rotman School of Management at the University of Toronto, Director, RiskLab Toronto and a principal of Sigma Asset Management. He may be reached at: seco@math.utoronto.ca

The first equation in (1) is the one given in Shadwick and Keating [2002]. The second is easy to derive from the first and may be more intuitive for investors. It provides a formulation where the benchmark can be taken to be a random variable, and gives the Omega as the ratio of the expected upside (outperformance) of the asset over the benchmark and its expected downside. Note that $\Omega(L) > 1$ (expected overperformance exceeds expected underperformance) exactly when $\mathbb{E}[R - L] > 0$.

Given Omega's growing popularity, constructing a portfolio that maximizes Omega is both a topical and practical problem. In this paper, we study the problem of finding a portfolio constructed from a set of instruments that maximizes the Omega given a benchmark L , which could be a number or a random variable. This optimization problem has also been considered by Avouyi-Dovi et al. (2004), who employ a global optimization approach, and Passow (2005), who considers a parametric approach using the four-parameter family of Johnson distributions.

At first sight, optimizing Omega is a difficult task, as the resulting optimization problem is not *convex*. We circumvent this difficulty by showing how a simple transformation of the problem variables (due to Charnes and Cooper [1962]) results in a linear program which is easy to solve on modern desktop computers. The transformation only works when the expected return of the optimal portfolio is larger than the benchmark (i.e. when the optimal Omega is greater than 1). We also discuss the alternatives available for optimizing the Omega when this condition fails. We provide an example that illustrates the methods by using market data to construct a fund of hedge funds that maximizes the Omega. Finally, we discuss how our formulation can be used to optimize Omega with a random benchmark (such as the return on an index), and illustrate its application to construct a portfolio of hedge funds that outperforms (in the sense of Omega) a given hedge fund index.

Performance measurement has always been a controversial topic, with the relative merits of various measures subject to continual debate. While this paper does not intend to add to this ongoing polemic, it is appropriate to review briefly some of the benefits and drawbacks of Omega. The key advantages of the Omega measure are its clear focus on the roles of upside and downside relative to a benchmark, and its flexibility. By assessing the Omega for different levels of the benchmark L , one can get a clear picture of the risk profile offered by the asset's return distribution (indeed, as pointed out by Shadwick and Keating [2002], knowing the value of $\Omega(L)$ for all L is equivalent to knowing the distribution of R). Further discussion of the advantages of Omega can be found in Shadwick and Keating [2002], and other papers

by the same authors.¹ Among the criticisms that have been leveled against Omega can be included the fact that Omega focuses solely on averages, while investors may be interested in other features of the returns distribution (for example, perhaps lower partial moment measures of risk more accurately reflect investors' aversion to large losses). Furthermore, how to choose a fixed benchmark level L to use for performance measurement is not obvious, and may have a significant impact on the resulting optimal portfolio. One way to mitigate this difficulty is by selecting a random benchmark L . We investigate this possibility within the context of Omega optimization later in this paper. Alternatively, the investor may choose to consider Omega optimal portfolios for a series of benchmarks L_i ; these portfolios may then serve as the building blocks in a more sophisticated portfolio construction process.

2 Omega Optimization

This section discusses different approaches for optimizing Omega. We formulate the problem in a market with a finite number of future scenarios, thus freeing ourselves from any of the particular distributional assumptions which often encumber performance measures. We first present a naive approach, which will serve as a performance benchmark. This approach formulates Omega optimization as a nonlinear constrained optimization problem, which may be solved using standard nonlinear programming (NLP) techniques. However, since the problem is nonconvex, it has the drawback that standard NLP solvers may return suboptimal portfolios. We introduce a transformation of variables that converts the nonlinear program into a linear programming problem. Achieving linearity requires dropping from the problem a particular set of constraints, referred to as complementarity constraints. This is possible since the complementarity constraints hold automatically for most financially relevant problems. In particular, they hold when the mean return of the optimal portfolio is larger than the benchmark (i.e. when the optimal Omega is greater than one). This is likely to occur in practice when the interest of the investor is risk management, in which case a low benchmark will be selected, or expected performance, in which case the instruments in the asset universe will tend to have high means.

As is typically the case when modeling investment returns, there are two possible ways to proceed: the parametric approach and the non-parametric one. Parametric approaches assume an underlying return distribution, whose parameters are estimated from historical data or calibrated to market prices. Non-parametric approaches forego such assumptions and instead use the his-

¹Available at www.financedevelopmentcentre.com.

toretical observations and the sample measure associated to them as the underlying portfolio distribution. Our approach follows the non-parametric model, which is also popular in stress-testing methodologies. The parametric approach to optimizing Omega discussed in Passow (2005) offers an interesting alternative to the approach taken in this paper.

We consider a collection of N assets from which an investor can construct a portfolio, and assume that there is a finite number S of possible future scenarios (which may come, for example, from a Monte-Carlo simulation of a continuous distribution function) with probabilities $p_i > 0$ for $i = 1, \dots, S$. The return of asset j under scenario i is denoted by R_{ij} . If an investor puts the fraction w_j of total wealth in asset j , then the portfolio return will be $(Rw)_i$ in scenario i , where R is the matrix with entries R_{ij} . Given a benchmark return L , equation (1) implies that the Omega of the portfolio with weights w may be written as:

$$\Omega(L) = \frac{\sum_{i=1}^S p_i u_i}{\sum_{i=1}^S p_i d_i}$$

where the positive and negative parts of the excess return over the benchmark (upside and downside) are given by:

$$\begin{aligned} (Rw)_i - L &= u_i - d_i & i = 1, \dots, S \\ u_i, d_i &\geq 0 & i = 1, \dots, S \end{aligned}$$

In order to have u and d truly represent the upside and downside it is necessary that $u_i \cdot d_i = 0$ for $i = 1, \dots, S$, to ensure that for each scenario either the upside or the downside is zero.²

Suppose that all the investor's wealth is to be invested in the asset portfolio, and that short selling is not permitted. This leads to the following restrictions on the portfolio weights:

$$\begin{aligned} \sum_{j=1}^N w_j &= 1 \\ w_j &\geq 0 \quad j = 1, \dots, N \end{aligned}$$

We allow the investor to put constraints on the portfolio, as long as they may be expressed as a system of linear inequalities, i.e. they have the form $Aw \leq b$. This encompasses a broad class of portfolio constraints, including position limits on individual assets and asset classes.

²This is necessary to avoid cases such as the following. Suppose $Rw_i = 5, L = 3$ so that $Rw_i - L = 2$. Without the constraint on u_i and d_i one could satisfy the equation with $u_i = 4$ and $d_i = 2$, in which case u_i and d_i would clearly not represent the upside and downside of the portfolio return under the scenario i .

2.1 The Straightforward Approach

The investor’s problem of finding the portfolio that optimizes the Omega is:

$$\Omega^*(L) = \max_{w,u,d} \frac{\sum_{i=1}^S p_i u_i}{\sum_{i=1}^S p_i d_i} \quad (2)$$

$$\sum_{j=1}^N R_{ij} w_j - u_i + d_i = L \quad i = 1, \dots, S \quad (3)$$

$$\sum_{j=1}^N w_j = 1$$

$$Aw \leq b$$

$$u_i, d_i, w_j \geq 0 \quad i = 1, \dots, S \quad j = 1, \dots, N$$

$$u_i \cdot d_i = 0 \quad i = 1, \dots, S \quad (4)$$

The above optimization problem is a nonconvex nonlinear program. It can be solved using specialized software for nonlinear programming, which looks for a portfolio satisfying a system of inequalities (the Kuhn-Tucker conditions) that are necessary (but not in general sufficient) for optimality. Such solvers can only find a “local” maximum, not necessarily the portfolio giving the true “global” solution (i.e. the portfolio having the largest $\Omega(L)$). A simple approach that often overcomes this difficulty is to call the solver with many different initial starting points, and select as the “true” optimal portfolio the best of the local optima computed from each of these starting points.

We solved this problem using historical scenarios on a set of 10 of the HFRX hedge fund indices. The index symbols, names, and moments of their returns distributions are given in Table (1) (the data and strategy definitions for the HFRX indices are available, after registration, from www.hedgefundresearch.com), using monthly return data from April 2003 to March 2006, for various values of the benchmark return L .³ The results are reported in Figure 1. The only trading constraint applied was the prohibition of short-selling. The optimal attainable Omega is compared at each level of the benchmark return with the Omega of a “strawman” portfolio with equal weights in each of the funds. We note that the figure does not compare the Omega curves (see Shadwick and Keating [2002]) of two different portfolios. Rather, at each value of L the “optimal Omega” portfolio is different, and its Omega is compared to that of the benchmark portfolio.⁴

³The choice of the empirical distribution on historical scenarios is made for illustrative purposes. The usual warning that the future will not be a repetition of the past applies. The identical techniques would be applied if instead one were to use Monte-Carlo scenarios based on a prescribed joint probability distribution for the returns.

⁴For the results in this figure, the nonlinear optimizer was run only once, with initial

2.2 Linear Programming Formulation

In this section, we introduce a transformation of variables (originally due to Charnes and Cooper [1962]) that makes the above nonconvex nonlinear program into a linear program, provided that we are willing to drop the complementarity constraint (4). We also discuss the consequences of omitting this constraint. Let:

$$t = \frac{1}{\sum_{i=1}^S p_i d_i}$$

and notice that if $\Omega^*(L) \in (0, \infty)$ then $t > 0$ is finite. Define the transformed variables:

$$\tilde{w}_j = w_j t \quad \tilde{u}_i = u_i t \quad \tilde{d}_i = d_i t \quad (5)$$

Observe that since $t > 0$, the nonnegativity of the transformed variables is equivalent to that of the original variables. Additionally, given the variables $t, \tilde{w}, \tilde{u}, \tilde{d}$ it is easy to return to the original variables using the inverse transformation:

$$w_j = \tilde{w}_j / t \quad u_i = \tilde{u}_i / t \quad d_i = \tilde{d}_i / t \quad (6)$$

Substituting the transformed variables into the problem (2) and dropping the constraint (4) results in the following optimization problem:

$$\begin{aligned} \tilde{\Omega}(L) &= \max_{\tilde{w}, \tilde{u}, \tilde{d}, t} \sum_{i=1}^S p_i \tilde{u}_i & (7) \\ \sum_{j=1}^N R_{ij} \tilde{w}_j - \tilde{u}_i + \tilde{d}_i - Lt &= 0 \quad i = 1, \dots, S \\ \sum_{j=1}^N \tilde{w}_j - t &= 0 \\ \sum_{i=1}^S p_i \tilde{d}_i &= 1 \\ A\tilde{w} - bt &\leq 0 \\ \tilde{u}_i, \tilde{d}_i, \tilde{w}_j &\geq 0 \quad i = 1, \dots, S \quad j = 1, \dots, N \end{aligned}$$

which is a linear program in the variables $\tilde{w}, \tilde{u}, \tilde{d}, t$.

Unfortunately, we cannot automatically assume that the optimal solution of (7) gives an optimal solution to (2) by reversing the transformation (5). This is because we have dropped the complementarity constraint (4) that $u_i \cdot d_i = 0$ under all scenarios. We are forced to make a more careful consideration of portfolio equal to the “strawman” portfolio.

the linear program (7). This leads to the following results, whose proofs are contained a technical appendix (available from the authors upon request).

1. The optimal value $\tilde{\Omega}(L)$ of the linear program satisfies $\tilde{\Omega}(L) \geq 1$. This can be understood as follows. Recall that (7) is a reformulation of (2) (through the change of variables (5) with the complementarity constraint (4) dropped). Without the complementarity constraint, both u and d can increase to infinity at the same rate, while maintaining a constant difference (satisfying (3)). The ratio in the objective function of (2) thus tends to one, and the optimal value of the linear program (7) (equivalently, of the nonlinear program (2)) without the constraint (4) must be greater than or equal to one. This argument does not imply that the optimal Omega $\Omega^*(L)$ is truly larger than one. As u and d increase without bound, they violate the complementarity constraint (4) and are therefore financially meaningless. *Only feasible solutions that satisfy the complementarity constraint are financially meaningful, as only in this case do u and d give the portfolio overperformance and underperformance, respectively.*
2. If $\tilde{\Omega}(L) > 1$, then the complementarity constraints $u_i \cdot d_i = 0$ are automatically satisfied, $\tilde{\Omega}(L) = \Omega^*(L)$, and the optimal solution of (2) may be obtained from the optimal solution to (7) by reversing the transformation (5). For example, if (u, d) is a feasible solution with objective value strictly greater than one, and violating the complementarity constraint by $u_i > d_i = \varepsilon > 0$, then the feasible solution (u^*, d^*) which is equal to (u, d) except that $u_i^* = u_i - \varepsilon$, $d_i^* = 0$ will have a strictly larger objective value (a similar argument works when $d_i > u_i$).

We solved the linear program (7) for the same hedge fund data used in the previous section. Figure 2 shows the optimal Omega computed with the linear program (7) compared to the results obtained using the nonlinear solver for the problem (2) from the previous section. The shape of the curve in Figure 2 is characteristic of such problems. It is clear that the Omega is a decreasing function of the threshold L , and this property carries over to the optimal value of Omega over a set of portfolios. Furthermore, using standard arguments from optimization, one can show that the optimal value function will always be convex in the region where $\tilde{\Omega}(L) > 1$. This reflects the fact that the rate at which the optimal Omega declines is a decreasing function of the threshold L (indeed, one should note that for L small enough, and with a finite number of scenarios, one will eventually have $\Omega(L) = \infty$, due to the truncation of the left tail; of course, this need not happen with continuous returns distributions). There are three key observations to make regarding this figure:

- The linear programming problem is easier to implement and allows a more efficient solution using standard software packages.
- There are points on the graph where the nonlinear solver has only identified a local, rather than a global solution. The linear solver always produces a global solution (given that $\tilde{\Omega}(L) > 1$).
- The transformation only works when $\tilde{\Omega}(L) > 1$.

3 The Case $\tilde{\Omega}(L) \leq 1$

The transformation of variables discussed above provides a fast and easy way to optimize $\Omega(L)$ for the financially important case $\tilde{\Omega}(L) > 1$. Unfortunately, when the optimal Omega is less than one, the linear program (7) won't work. This is because it will produce a solution with optimal value 1 that violates the complementarity constraints $u_i \cdot d_i = 0$ and is therefore financially meaningless. In the case where the optimal Omega is less than one, the investor wishing to solve problem (2) has three alternatives:

Nonlinear Programming

One may attempt to apply nonlinear programming techniques to try to produce a solution to (2). For example, one could solve the problem with standard NLP software and many initial points, taking as the optimal portfolio the one giving the largest local maximum returned by the solver.

The advantages of this approach are that it is still relatively efficient (in terms of computation time and storage) and can be implemented easily on a standard desktop computer. The drawback is that one cannot be assured that the portfolio returned by the optimizer is truly "optimal".

Global Optimization

Another alternative is to employ more advanced global optimization algorithms or optimization heuristics that may be tailored to the problem at hand, such as Tabu search (see Glover [2005]), or Threshold Accepting (see Avouyi-Dovi et al. [2004]).

The advantage of this approach is that it is much more likely to produce a true optimum than the above, more naive, strategy. Its drawbacks include the difficulty in verifying optimality, and expertise required for its implementation.

Integer Programming

This method adds the following constraints to the linear program (7):

$$\begin{aligned} u_i &\leq Mz_i \\ d_i &\leq M(1 - z_i) \end{aligned} \tag{8}$$

where z_i are binary variables, i.e. $z_i = 0$ or 1 and M is a large number. It is easy to see that it is impossible to satisfy the constraints (8) and have both $u_i > 0$ and $d_i > 0$. Specifically, if $z_i = 1$ then $d_i = 0$, while if $z_i = 0$, then $u_i = 0$. Thus augmenting the linear program (7) with the constraints (8) effectively enforces the complementarity constraints (4). One can then use standard integer (linear) programming techniques to solve the Omega optimization problem.

The advantage of this method is that it is straightforward to implement. The disadvantage is that the computational effort required to find the optimal portfolio can increase dramatically. The method introduces a new binary variable z_i for each scenario. Thus for small scenario sets (as often found in the hedge fund industry), the method will be relatively efficient. However, as the number of scenarios grows, the method will become increasingly cumbersome, and the time taken to solve the problem will grow much faster than that required by any of the other methods.

4 Random Benchmark Optimization

This section considers the optimization of Omega with a random benchmark. Recall from equation (1) that a simple argument allows one to transform the standard definition of Omega given by Shadwick and Keating [2002], into the following:

$$\Omega(L) = \frac{E[\max(R - L, 0)]}{E[\max(L - R, 0)]}$$

Observe that this definition is valid whether L is a constant or a random variable (in which case one needs the joint distribution of R and L in order to calculate Omega)⁵. Applying this simple observation to the optimization problems discussed in the previous section, we see that these problems remain exactly the same except that in each formulation the constant L must be replaced by the scenario dependent parameter L_i .

⁵We note here that when L is random the first formulation in equation (1) no longer holds directly. Nonetheless, with a minor modification we still have $\Omega(L) = \frac{\int_0^\infty (1 - G(x)) dx}{\int_{-\infty}^0 G(x) dx}$ where G is the distribution function of the random variable $R - L$.

We solve the optimization problem with a random benchmark to construct portfolios of hedge fund indices. Specifically, we consider the same set of HFRX indices used in the previous section, and solve (7) for two different benchmarks. The first benchmark is the Salomon-Smith-Barney (SSB) Government Bond Index and the second is the HFRI Fund of Funds Composite Index. For simplicity, we again consider the empirical distribution for our scenario set, placing equal weight on each historically observed monthly return.

Solving (7) for the SSB Government Bond Index yields a portfolio with an optimal Omega of $\Omega^*(L) = 5.2644$, consisting of about 41% in each of the Distressed Securities (0.4122) and Event-Driven (0.4062) indices and 18% (0.1816) in the Equity Market Neutral Index. (recall that now L represents the *random* return on the SSB Government index). The HFRI benchmark is more difficult to beat; the optimal Omega is only 1.0356. Furthermore, the optimal portfolio is not diversified, as it invests only in the event driven index. Figure 3 plots histograms of returns for each of the optimal portfolios. Table 2 gives summary statistics for each of the distributions. In this case, the dramatically more conservative SSB benchmark actually produces only a slightly less conservative optimal portfolio, by traditional measures. As is the case whenever performance is measured against a benchmark, the influence of the choice of benchmark is critical, and its impact is worthy of further investigation.

We also use the random benchmark data to test the stability of the optimal portfolio through time. We re-solved the linear program (7) using the SSB Government Bond Index as benchmark, with different sets of data. Specifically, we envisioned a manager using data beginning in April 2003, and re-solving the optimal portfolio problem each month, beginning with April 2005, and continuing until March 2006, adding the new data to the optimal portfolio problem as it became available. The optimal Omegas and optimal portfolio compositions for this investor are shown in Figure (4). Of particular significance is the fact that the stability of the optimal portfolio (and indeed of the Omegas of individual instruments as well) depends crucially on the stability of the underlying data.

5 Conclusion

The Omega has recently enjoyed much popularity as a performance measure for analyzing the returns on alternative assets. In response to the challenge of optimizing Omega, this paper presents several theoretical results, together with their implications for computational implementations. In particular, we have demonstrated that when the mean of the optimal portfolio is greater

than that of the benchmark, a simple transformation of variables allows this problem to be solved using linear programming. We also discussed the alternatives available when the condition on the mean of the optimal portfolio fails. Furthermore, we have demonstrated that our approach applies to the maximization of Omega with a single fixed benchmark rate, or with a benchmark random variable (such as the return on an index), and is easy to implement using standard software on modern desktop computers.

References

- [1] S. AVOUYI-DOVI, A. MORIN, AND D. NETO, *Optimal asset allocation with Omega function*, tech. report, Banque de France, 2004. Research Paper.
- [2] A. CHARNES AND W. COOPER, *Programming with linear fractional functionals*, Naval Research Logistics Quarterly, 9 (1962), pp. 181–186.
- [3] R. DEMBO AND H. MAUSSER, *The put/call efficient frontier*, Algo Research Quarterly, 3 (2000), pp. 13–26.
- [4] R. DEMBO AND D. ROSEN, *The practice of portfolio replication: A practical overview of forward and inverse problems*, Annals of Operations Research, 85 (1999), pp. 267–284.
- [5] F. GLOVER, *A parametric (scaled penalty) ts method for a linear complementarity problem*, tech. report, University of Colorado, 2005.
- [6] P. KAPLAN AND J. KNOWLES, *A generalized downside-risk performance measure*, tech. report, Morningstar, 2003. Research Paper.
- [7] A. PASSOW, *Omega portfolio construction with Johnson distributions*, Risk, (2005), pp. 85–90.
- [8] W. SHADWICK AND C. KEATING, *A universal performance measure*, Journal of Portfolio Measurement, Spring (2002).
- [9] F. SORTINO, R. VAN DER MEER, AND A. PLANTINGA, *The Dutch triangle*, Journal of Portfolio Management, 26 (1999), pp. 50–58.

Symbol	Strategy	Mean	StDev	Skewness	Excess Kurtosis
HFRXGL	Global Hedge Fund	0.5419	1.1120	-0.3150	-0.3589
HFRXEW	Equal Weighted Strategies	0.4242	0.8838	-0.2720	-0.4177
HFRXCA	Convertible Arbitrage	-0.0159	1.3414	-0.4912	-0.0129
HFRXDS	Distressed Securities	0.6991	1.1915	0.0194	-0.6752
HFRXEH	Equity Hedge	0.6443	1.5185	-0.1739	-0.1775
HFRXEMN	Equity Market Neutral	-0.0389	0.7188	-0.3896	0.4481
HFRXED	Event-Driven	0.8236	1.3392	-0.2544	-0.4211
HFRXM	Macro Index	0.5848	1.8489	-0.0246	-0.5075
HFRXMA	Merger Arbitrage	0.4047	0.9660	-1.0727	2.1659
HFRXRVA	Relative Value Arbitrage	0.2769	0.6696	-0.2245	-0.5282

Table 1: Summary Statistics for HFRX Tradeable Indices.




Figure 1: Optimal Omegas using the Naive Approach

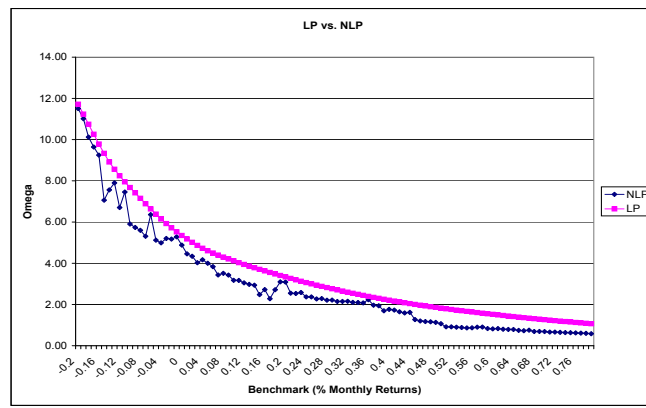


Figure 2: Comparison of Optimal Omegas using NLP and LP

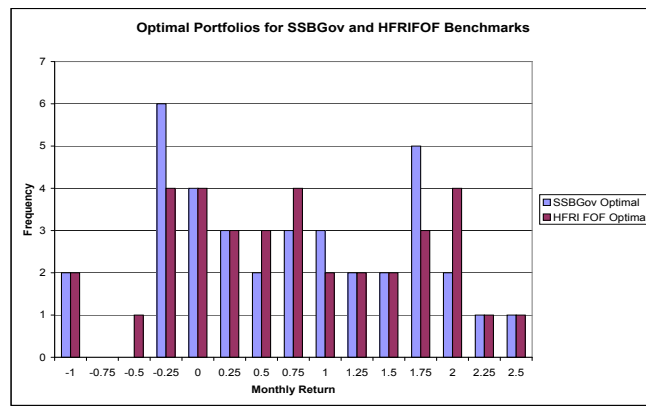


Figure 3: Optimal Portfolios with Different Benchmarks

Portfolio	Mean	StDev	Skewness	Excess Kurtosis
SSBGov Benchmark	0.0019	0.0149	-0.8237	1.4609
SSBGov Optimal	0.6156	0.9564	-0.0445	-0.7960
SSBGov Excess	0.6138	0.9606	-0.548	-0.7861
HFRIFOF Benchmark	0.8025	1.0493	-0.3871	-0.3303
HFRIFOF Optimal	0.6098	0.9521	-0.0088	-0.8420
HFRIFOF Excess	-0.1927	1.4103	0.2491	-0.3372

Table 2: Descriptive Statistics of Return Distributions.

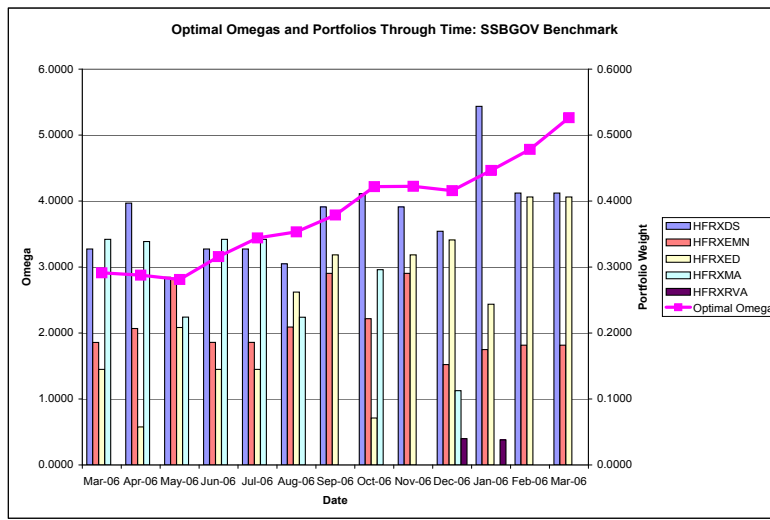


Figure 4: Optimal Portfolios and Portfolio Weights through Time

A Proofs of Results

Proposition 1. *The optimal value $\tilde{\Omega}(L)$ of the linear program satisfies $\tilde{\Omega}(L) \geq 1$. In the case that $\tilde{\Omega}(L) > 1$, then the complementarity constraints $u_i \cdot d_i$ are automatically satisfied, $\tilde{\Omega}(L) = \Omega^*(L)$, and the optimal solution of (2) may be obtained from the optimal solution to (7) by applying the inverse transformation (6).*

Proof. Substituting the inverse transformation (6) to the linear program for $\tilde{\Omega}(L)$ yields the optimization problem (2) with the complementarity constraints missing. More specifically, the optimization problem becomes:

$$\begin{aligned}
 \tilde{\Omega}(L) &= \max_{w,u,d} \frac{\sum_{i=1}^S p_i u_i}{\sum_{i=1}^S p_i d_i} \\
 &\sum_{j=1}^N w_j = 1 \\
 &\sum_{j=1}^N r_{ij} w_j - u_i + d_i = L \quad i = 1, \dots, S \\
 &Aw \leq b \\
 &u_i, d_i \geq 0 \quad i = 1, \dots, S \\
 &w_j \geq 0 \quad j = 1, \dots, N
 \end{aligned} \tag{9}$$

Let (w, u, d) be a feasible solution to the above problem, and let $\mathbf{1}$ be a vector of ones. Then for any $m > 0$, $(w, u + m\mathbf{1}, d + m\mathbf{1})$ is another feasible solution and

$$\lim_{m \rightarrow \infty} \frac{\sum_{i=1}^S p_i (u_i + m)}{\sum_{i=1}^S p_i (d_i + m)} = 1$$

and so $\tilde{\Omega}(L) \geq 1$.

Now suppose that $\tilde{\Omega}(L) > 1$. We must show that the complementarity slackness conditions $u_i \cdot d_i = 0$ hold automatically. Suppose to the contrary that for some i , we have $u_i > 0$ and $d_i > 0$, where (w, u, d) is an optimal solution. Let $w^* = w$, $\varepsilon = \min\left(\frac{u_i}{2}, \frac{d_i}{2}\right)$, $u^* = u - \varepsilon e_i$, $d^* = d - \varepsilon e_i$ where e_i is a vector with 1 in the i th position and 0 elsewhere. Then it is clear that (w^*, u^*, d^*) is feasible and

$$\begin{aligned}
 p^T u &> p^T u^* = p^T u - p_i \varepsilon > 0 \\
 p^T d &> p^T d^* = p^T d - p_i \varepsilon > 0
 \end{aligned}$$

so that

$$\frac{p^T u^*}{p^T d^*} > \frac{p^T u}{p^T d}$$

follows from the elementary fact that $a, b, c > 0$ with $a > b > c$ implies that $\frac{a-c}{b-c} > \frac{a}{b}$ ($a > b$ because $\tilde{\Omega}(L) > 1$). \square