Optimising omega

Optimising a portfolio's omega generally requires non-linear optimisation methods. ***Helmut Mausser*****,** ***David Saunders*** *and* ***Luis Seco*** *show that, under suitable conditions, a simple change of variables transforms the problem into a linear program that is much more readily solved*

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: they reflect the market practice of assessing performance against a benchmark, and they account for the asymmetry in returns distributions by separately considering upside and downside. Omega is a recent measure introduced by Shadwick & Keating (2002) possessing both of these characteristics. It is defined, for an asset with return R (a random variable) and benchmark return L as:

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

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 $D_{\lambda}(L) = E[\max(R - L, 0)] - \lambda E[\max(L - R, 0)]$ studied by Dembo and his colleagues (for example, Dembo & Rosen, 1999, and Dembo & Mausser, 2000) and the kappa ratios:

$$\kappa_n(L) = E[R-L] / \sqrt[n]{E\left[\max(L-R,0)^n\right]}$$

introduced by Kaplan & Knowles (2003) as generalisations of the Sortino ratio (see, for example, Sortino, Van der Meer & Plantinga, 1999).

The first equation in (1) is the one given in Shadwick & 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 omega as the ratio of the expected upside (over-performance) of the asset over the benchmark and its expected downside. Note that $\Omega(L) > 1$ (expected over-performance exceeds expected under-performance) exactly when $E[R - L] > 0$.

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

At first sight, optimising omega is a difficult task, as the resulting optimisation problem is not convex. We circumvent this difficulty by showing how a simple transformation of the problem variables (due to Charnes & Cooper, 1962) results in a linear program that 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 (that is, when the optimal omega is greater than one). We also discuss the alternatives available for optimising 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 maximises omega. Finally, we discuss how our formulation can be used to optimise 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 article 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 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 & 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 & Keating (2002), and other papers by the same authors.[1] Among the criticisms that have been levelled 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 optimisation later in this article. 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.

Omega optimisation

This section discusses different approaches for optimising omega. We formulate the problem in a market with a finite number of

---

[1] *Available at* www.financedevelopmentcentre.com

A. Summary statistics for HFRX tradable indexes

Symbol	Strategy	Mean	Standard deviation	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

future scenarios, thus freeing ourselves from any of the particular distributional assumptions that often encumber performance measures. We first present a naive approach, which will serve as a performance benchmark. This approach formulates omega optimisation as a non-linear constrained optimisation problem, which may be solved using standard non-linear programming techniques. However, since the problem is non-convex, it has the drawback that standard non-linear programming solvers may return sub-optimal portfolios. We introduce a transformation of variables that converts the non-linear 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 (that is, 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 modelling 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 forgo such assumptions and instead use the historical observations and the sample measure associated with 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 optimising omega discussed in Passow (2005) offers an interesting alternative to the approach taken in this article.

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:

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

To have u and d truly represent the upside and downside it is necessary that $u_i d_i = 0$ 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 expressesd as a system of linear inequalities, that is, they have the form $Aw \leq b$. This encompasses a broad class of portfolio constraints, including position limits on individual assets and asset classes.

The straightforward approach. The investor's problem of finding the portfolio that optimises the omega is:

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

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

$$\begin{aligned} \sum_{j=1}^{N} w_j = 1 \\ Aw \leq b \\ u_i, d_i, w_j \geq 0 \quad i = 1,\dots,S \quad j = 1,\dots,N \\ u_i d_i = 0 \quad i = 1,\dots,S \end{aligned} \tag{4}$$

The above optimisation problem is a non-convex non-linear program. It can be solved using specialised software for non-linear 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

---

[2] *This 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*

1 Optimal omegas using the naive approach

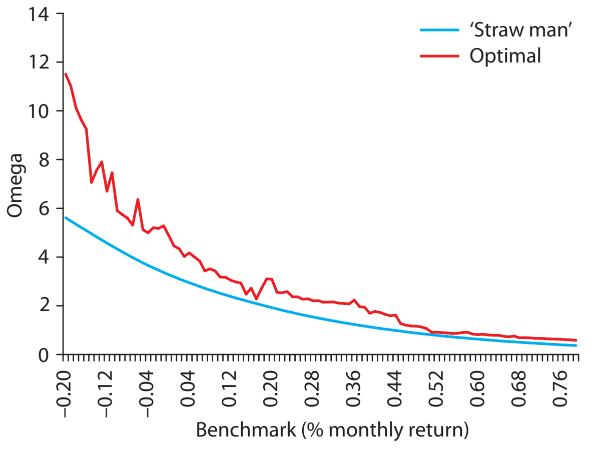

'global' solution (that is, the portfolio having the largest $\Omega(L)$). A simple approach that often overcomes this difficulty is to call the solver with many different initial starting points, and select as the 'true' optimal portfolio the best of the local optima computed from each of these starting points.

We solved this problem using historical scenarios on a set of 10 of the HFRX hedge fund indexes. The index symbols, names and moments of their returns distributions are given in table A (the data and strategy definitions for the HFRX indexes are available, after registration, at *www.hedgefundresearch.com*), using monthly return data from April 2003–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 'straw man' portfolio with equal weights in each of the funds. We note that the figure does not compare the omega curves (see Shadwick & Keating, 2002) of two different portfolios. Rather, at each value of L the 'optimal omega' portfolio is different, and its omega is compared with that of the benchmark portfolio.[4]

■ **Linear programming formulation.** In this section, we introduce a transformation of variables (originally due to Charnes & Cooper, 1962) that makes the above non-convex non-linear 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:

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

Observe that since $t > 0$, the non-negativity 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:

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

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

$$\tilde{\Omega}(L) = \max_{\tilde{w},\tilde{u},\tilde{d},t} \sum_{i=1}^{S} p_i \tilde{u}_i$$

$$\begin{gathered}
\sum_{j=1}^{N} R_{ij}\tilde{w}_j - \tilde{u}_i + \tilde{d}_i - Lt = 0 \qquad i = 1,\dots,S \\
\sum_{j=1}^{N} \tilde{w}_j - t = 0 \\
\sum_{i=1}^{S} p_i \tilde{d}_i = 1 \\
A\tilde{w} - bt \le 0 \\
\tilde{u}_i, \tilde{d}_i, \tilde{w}_j \ge 0 \qquad i = 1,\dots,S \qquad j = 1,\dots,N
\end{gathered} \tag{7}$$

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 $u_i d_i = 0$ under all scenarios. We are forced to make a more careful consideration of the linear program (7). This leads to the following results, whose proofs are contained in a technical appendix (available from the authors upon request):

■ 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 non-linear 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 over-performance and under-performance, respectively.

■ If $\tilde{\Omega}(L) > 1$, the complementarity constraints $u_i d_i = 0$ are automatically satisfied, $\tilde{\Omega}(L) = \Omega^*(\mathrm{L})$, 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 with the results obtained using the non-linear 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 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

---

[3] *The 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 one were to use Monte Carlo scenarios based on a prescribed joint probability distribution for the returns*

[4] *For the results in this figure, the non-linear optimiser was run only once, with initial portfolio equal to the 'straw man' portfolio*

standard arguments from optimisation, 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 small enough L, 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 non-linear 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 $\tilde{\Omega}(L) > 1$).

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

The transformation of variables discussed above provides a fast and easy way to optimise $\Omega(L)$ 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 one that violates the complementarity constraints $u_i d_i = 0$ 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:

- *Non-linear programming.* One may attempt to apply non-linear programming techniques to try to produce a solution to (2). For example, one could solve the problem with standard non-linear programming 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 optimiser is truly 'optimal'.

- *Global optimisation.* Another alternative is to employ more advanced global optimisation algorithms or optimisation heuristics that may be tailored to the problem at hand, such as Tabu search (see Glover, 2005) or threshold accepting (see Avouyi-Dovi, Morin & Neto, 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):

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

where z_i are binary variables, that is, $z_i = 0$ 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 optimisation 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.

2 Comparison of optimal omegas using non-linear programming and linear programming

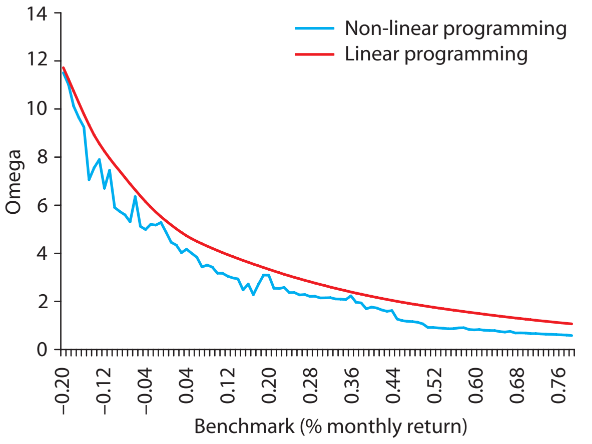

Thus for small scenario sets (as are often found in the hedge fund industry), the method will be relatively efficient. However, as the number of scenarios grows, the method will become more cumbersome, and the time taken to solve the problem will grow much faster than that required by any of the other methods.

Random benchmark optimisation

This section considers the optimisation 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 & Keating (2002), into the following:

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

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 to calculate omega).[5] Applying this simple observation to the optimisation 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.

We solve the optimisation problem with a random benchmark to construct portfolios of hedge fund indexes. Specifically, we consider the same set of HFRX indexes 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) indexes and 18% (0.1816) in the Equity Market Neutral

---

[5] *We 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:*

$$\Omega(L) = \frac{\int_0^\infty \left(1 - G(x)\right) dx}{\int_{-\infty}^0 G(x)\, dx}$$

where G is the distribution function of the random variable R − L

3 Optimal portfolios with different benchmarks

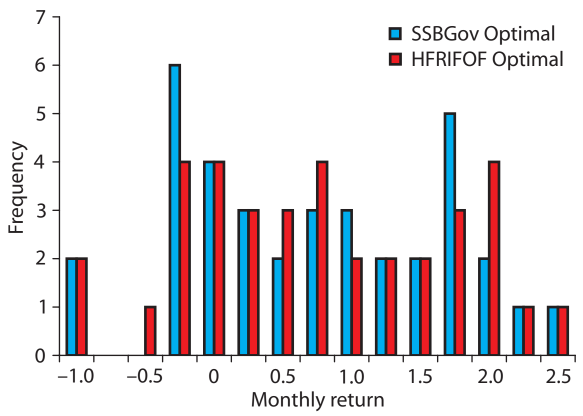

4 Optimal portfolios and portfolio weights through time

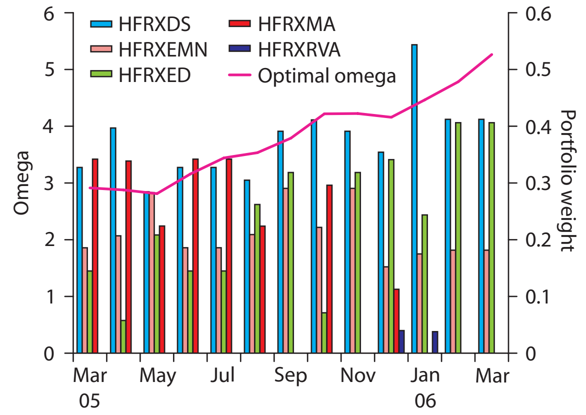

B. Descriptive statistics of return distributions

Portfolio	Mean	Std dev	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

Index (recall that now L represents the random return on the SSB Goverment 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 B 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 a 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.

Conclusion

Omega has recently enjoyed much popularity as a performance measure for analysing the returns on alternative assets. In response to the challenge of optimising omega, this article 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 maximisation 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. ■

Helmut Mausser is a senior mathematician at Algorithmics. David Saunders is an assistant professor in the Department of Statistics and Actuarial Science at the University of Waterloo. Luis Seco is a professor in the Department of Mathematics and the Rotman School of Management at the University of Toronto, a director at RiskLab Toronto and a principal of Sigma Analysis & Management. The authors are grateful to Fred Glover, Gustavo Comezaña and two anonymous referees for helpful comments and suggestions. Email: hmausser@algorithmics.com, dsaunders@math.uwaterloo.ca, seco@math.utoronto.ca

References

Avouyi-Dovi S, A Morin and D Neto, 2004
Optimal asset allocation with omega function
Technical report, Banque de France

Charnes A and W Cooper, 1962
Programming with linear fractional functionals
Naval Research Logistics Quarterly 9, pages 181–186

Dembo R and H Mausser, 2000
The put/call efficient frontier
Algo Research Quarterly 3, pages 13–26

Dembo R and D Rosen, 1999
The practice of portfolio replication: a practical overview of forward and inverse problems
Annals of Operations Research 85, pages 267–284

Glover F, 2005
A parametric (scaled penalty) ts method for a linear complementarity problem
Technical report, University of Colorado

Kaplan P and J Knowles, 2003
A generalized downside-risk performance measure
Technical report, Morningstar

Passow A, 2005
Omega portfolio construction with Johnson distributions
Risk April, pages 85–90

Shadwick W and C Keating, 2002
A universal performance measure
Journal of Performance Measurement, spring, pages 59–84

Sortino F, R van der Meer and A Plantinga, 1999
The Dutch triangle
Journal of Portfolio Management 26, pages 50–58