Using Lagrangian Relaxation to Obtain Small Portfolios

LEONID KOPMAN, SHUCHENG SCOTT LIU, AND DONG SHAW

LEONID KOPMAN
is a senior associate of research at MSCI Barra in Berkeley, CA.
leonid.kopman@mscibarra.com

SHUCHENG SCOTT LIU
is a vice president of research at MSCI Barra in Berkeley, CA.
scott.liu@mscibarra.com

DONG SHAW
is the founder and president of Analytic Portfolio LLC in New York, NY.
dong.shaw@analyticportfolio.com

Portfolio managers using mean-variance optimization for portfolio construction often want to impose a limit on the number of stocks in their portfolios. A common practice for passive fund managers is to purchase only a subset of the issues in an index. Active managers also often place a limit on the number of issues in their portfolios. As a result, portfolio optimization problems with a limit on the number of assets are important in practice. Unfortunately, putting a limit on the number of assets can make a typical portfolio optimization problem very difficult to solve.

A practitioner can approach the problem in various ways. The most common approach is to construct an approximate solution instead of attempting to solve the problem to optimality. A number of heuristics for this problem have been developed over the years. One of the latest methods is the diversity heuristic (DH), described in Jansen and van Dijk [2002]. In recent years, because of dramatically enhanced computer power and greatly improved computational techniques, it has become possible to directly tackle the exact problem. Charpin and Lacaze [2007] advocated using binary variables to formulate portfolio optimization problems with a maximum number of assets and/or minimum buy-in thresholds. They solved these problems to optimality with quadratic programming-based branch-and-bound. To justify their choice, they compared optimal portfolios obtained using their method (which they call the binary variable method) with suboptimal portfolios obtained using Jansen and van Dijk's diversity heuristic.

A provably optimal solution is, of course, always preferable, in terms of quality, to a suboptimal solution. Unfortunately, it is not always possible to compute the optimal solution within a reasonable time, even with modern computing power. The largest problems solved by Charpin and Lacaze have 250 assets in the selection universe; computation times are not reported. Larger asset selection universes will, undoubtedly, present computational challenges. Running branch-and-bound with a timeout has its shortcomings as well in that a method to obtain optimal solutions reduces to a heuristic and the quality of solutions provided may leave much to be desired.

We propose a new method for solving these problems. It combines the best attributes of both approaches just referenced. Our approach can solve smaller problems to optimality in reasonable time. For larger problems, our approach produces heuristic solutions with estimates of their proximity to optimality. In the case of passive index tracking, for instance, this estimate is the distance between the heuristic solution's tracking error (TE) and the bound on the TE of the optimal portfolio. The TEs of our heuristic solutions are typically lower than the TEs of portfolios obtained with DH or running branch-and-bound based on binary variables with timeout.

WHICH COVARIANCE MATRIX TO USE?

In a portfolio optimization problem, risk is typically taken into account via the asset return covariance matrix. If h is a $n \times 1$ vector of portfolio weights, and Q is an $n \times n$ matrix of asset return covariances, the term representing risk in the problem is $h'Qh$ for portfolio risk or $(h - h_b)'Q(h - h_b)$ for tracking error, where h_b is a $n \times 1$ vector of benchmark weights. This term can be present in the objective function with a suitable risk aversion coefficient. Asset return covariance matrix Q can be computed statistically using asset return histories. This method was utilized by Jansen and van Dijk [2002], as well as by Charpin and Lacaze [2007]. A more standard practice in modern portfolio optimization, however, is to use a multiple-factor model (MFM). MFMs have been successfully used since the 1970s in commercial portfolio optimization tools. The basic premise of an MFM is that many influences act on the volatility of a stock, and these influences (factors) are often common across many stocks. According to Rosenberg [1974], because the factors can represent the component of return seen by financial analysts, the multiple-factor model is a natural representation of the real environment.

In an MFM, the return of each security is decomposed into a part attributed to common factors and an idiosyncratic part called the asset-specific return. This can be written mathematically as $r = Xf + u$, where r is the vector of total asset excess returns over the risk-free rate, f is the vector of factor returns, X is the matrix of linear coefficients representing the sensitivity of assets to the factors (factor-loading matrix), and u is the vector of the asset-specific (non-factor) returns. Since the factors and asset-specific risks are uncorrelated, the asset return covariance matrix Q can be written as $Q = XFX' + D$, where D is a diagonal matrix of specific risk variance, and F is the covariance matrix of factors.

Our method for solving portfolio optimization problems with a limit on the number of assets, presented in the next section, utilizes the special structure of an MFM covariance matrix. Due to the widespread use of MFMs by financial practitioners, we believe that the reliance of our method on MFM does not present a real-world limitation.

LAGRANGIAN RELAXATION METHOD

Our approach to finding small portfolios is based on a new procedure that finds the bound on the value of the utility function achieved by the optimal small portfolio, and, in the process, also produces good heuristic solutions. The procedure is applicable to both small problems (in a branch-and-bound framework) and large problems (by itself, producing a good solution accompanied by the bound it finds). It utilizes the special structure of the asset return covariance matrix yielded by the MFM and is based on the Lagrangian relaxation.

The bounding method involves splitting the original problem into two problems that are connected to each other with auxiliary (dual) variables. Both resulting problems are easy to solve. One is a portfolio optimization problem without the limit on the number of assets, and the other is solved simply by ordering assets according to their "attractiveness." The upper bound is given by the sum of the two problems' objective functions. The best upper bound is obtained by iterating on the dual variables. The ordering of the assets obtained in the process of solving the second problem gives rise to a heuristic—if a portfolio with a maximum of K assets is desired, restrict the original problem to K "best" assets from the ordering and solve the modified problem to obtain optimal weights. A more detailed description follows.[1]

To simplify the exposition, we work with a bare-bones problem; that is, the objective function contains return and variance terms, and constraints consist of the holding constraint and bounds on asset holdings. We denote the constraint limiting the number of assets in the portfolio as $nnz(h) \leq K$ (i.e., number of nonzeros in h is limited by K). With this constraint added, the portfolio optimization problem has the form[2]

$$\begin{aligned} &\max c'h - \lambda h'(XFX' + D)h \\ &\sum h_i = 1 \\ &l \leq h \leq u \\ &nnz(h) \leq K \end{aligned} \tag{1}$$

Problem (1) can be rewritten as

$$
\begin{aligned}
f^* = \max c'z - \lambda h'XFX'h - \lambda z'Dz \\
\sum h_i = 1 \\
h = z \\
l \le h \le u, l \le z \le u \\
nnz(z) \le K
\end{aligned} \tag{2}
$$

To this point, we have not changed much. We introduced variables z that mirror h, and we substituted z for h in two places. Therefore, the optimum solution of Problem (2) is the same as the optimum solution of Problem (1). Then, we introduce the dual vector μ and obtain a Lagrangian relaxation with respect to constraints h = z. The easiest way to think about our approach is that, instead of imposing $h = z$ in a strict way, we instead penalize deviation between h and z, with μ being the penalty parameter. The relaxed problem becomes

$$
\begin{aligned}
\max c'z - \lambda h'XFX'h - \lambda z'Dz - \mu'(z - h) \\
\sum h_i = 1 \\
l \le h \le u, l \le z \le u \\
nnz(z) \le K
\end{aligned} \tag{3}
$$

Now, notice that z variables are no longer connected by constraints to other variables in Problem (3). Indeed, for a given μ, optimal values of z can be determined independently by solving a separate problem. Hence, Problem (3) splits into the two following problems:

$$
\begin{aligned}
L^1(\mu) = \max(c - \mu)'z - \lambda z'Dz \\
nnz(z) \le K \\
l \le z \le u
\end{aligned} \tag{4}
$$

and

$$
\begin{aligned}
L^2(\mu) = \max \mu'h - \lambda h'XFX'h \\
\sum h_i = 1 \\
l \le h \le u
\end{aligned} \tag{5}
$$

Owing to the relaxation (i.e., not imposing the constraint $h = z$ in Problem (3)), $f^* \le L^1(\mu) + L^2(\mu)$ for any μ. Note that Problem (5) is a portfolio optimization problem without the limit on the number of assets and can be solved quickly by using the usual methods. Meanwhile, Problem (4) is even simpler. It can be solved by computing the "best possible contribution" from each variable, ordering the asset-holding variables from best to worst, and then selecting K best; the method could be characterized as a type of greedy algorithm. To summarize, the procedure to obtain an upper bound on the optimal objective-function value of the portfolio optimization problem with a limit on the number of assets consists of solving Problems (4) and (5), and adding their optimal objective function values.

Note that $f^* \le L^1(\mu) + L^2(\mu)$ for all values of μ. We find the μ that provides the best upper bound on f^* by iterating on μ using a *subgradient algorithm*.[3] At every iteration of the algorithm, a byproduct of solving Problem (4) is an ordering of holding variables with respect to their attractiveness and a set of K most attractive assets. An obvious byproduct heuristic for solving the original problem consists of restricting the problem to just those K assets and then solving it.

The bounding procedure and the byproduct heuristic were implemented within the branch-and-bound framework. This Lagrangian relaxation–based branch-and-bound is different from the branch-and-bound based on QP formulation with binary variables, advocated by Charpin and Lacaze [2007]. Our approach provides tighter bounds on the optimal value of the objective function. Hence, as our computational experience shows, we can find optimal portfolios for problems with small- to medium-sized selection universes faster. Larger problems typically cannot be solved to optimality in reasonable time. However, after the bounding procedure is run on the original problem and before any branching is done, we find a solution accompanied with a bound on the optimal value, with the gap between them typically small (equivalent to a few basis points of tracking error).

Similar to the binary variable method (and unlike the diversity heuristic), the Lagrangian relaxation method can be extended to accommodate minimum holding thresholds, minimum transaction thresholds, and so forth.

NUMERICAL RESULTS

We used several portfolio optimization problems to compare the Lagrangian relaxation method with the binary variable–based branch-and-bound and the diversity heuristic approaches. Asset covariance matrices were given by the Barra USE3L risk model for January 31, 2008. We used the following indices: S&P100, FT100, S&P500, FR2000, and FR3000. We constructed a passive (index tracking) and an active (with randomly generated asset returns) long-only portfolio optimization problem for each index. The maximum number of assets was selected to be 20 for 100-asset problems, 50 for 500-asset problems, and 75 for 2000- and 3000-asset problems.

In addition to running the Lagrangian relaxation method, we formulated these problems as mixed-integer quadratic programming problems with binary variables. We attempted to solve those with ILOG CPLEX 10.0—a multi-purpose commercial optimization package. Branch-and-bound was run on each test case for 600 seconds. We also implemented the diversity heuristic with parameters described by Jansen and van Dijk [2002] and ran it on our test problems.

Exhibit 1 shows our results for the index tracking problems. We ran Lagrangian relaxation–based branch-and-bound on 100-asset problems and solved them to optimality. For the rest of the problems, Lagrangian relaxation procedure was run on the original problem, and no branching was done. In one of five cases, the binary variable method found the optimal solution, but its optimality was not proven. In fact, none of the test cases was solved to optimality with this method. Exhibit 1 reports the best solution found with each method, as well as the global bound reported at the end of the run (converted into TE terms), where applicable.

Exhibit 2 presents the results for the active problems (note these are maximization problems). The 100-asset and 500-asset cases were solved to optimality with our method. For the remaining cases, the Lagrangian

EXHIBIT 1
Computational Results for Index Tracking Problems

Test Case	Lagrangian Relaxation			Binary Variable Method			Diversity Heuristic
	TE, %	Bound, %	Time, s	TE, %	Bound, %	Time, s	TE, %
sap100	2.474	OPT	132	2.474	0.95	600	2.946
ft100	2.019	OPT	50	2.022	1.03	600	2.276
sap500	1.874	1.858	18	1.906	0.12	600	2.380
fr2000	3.383	3.374	51	3.709	0.01	600	4.145
fr3000	1.662	1.649	100	2.872	0.01	600	2.736

EXHIBIT 2
Computational Results for Active Problems

Test Case	Lagrangian Relaxation			Binary Variable Method			Diversity Heuristic
	Obj	Bound	Time, s	Obj	Bound	Time, s	Obj
sap100a	3.25897	OPT	26	3.25897	3.35506	600	3.2375
ft100a	2.89466	OPT	25	2.89466	2.95667	600	2.7780
sap500a	3.98442	OPT	268	3.98442	4.09022	600	3.9487
fr2000a	3.60332	3.60498	36	3.54462	4.22722	600	2.6819
fr3000a	3.88631	3.88746	68	3.87993	4.03686	600	3.6526

relaxation procedure was run on the original problem, and no branching was done. In three of five cases, the binary variable method found the optimal solution, but its optimality was not proven. Again, none of the test cases was solved to optimality with this method. Exhibit 2 shows the best solution value (obj) found with each method, as well as the global upper bound reported at the end of the run, where applicable.

The results presented in this article indicate that the Lagrangian relaxation method is faster and exhibits better performance than the binary variable method. The diversity heuristic appears to be inferior to both methods, which supports the findings of Charpin and Lacaze [2007].

CONCLUSION

Utilizing the special structure of the asset covariance matrix that is yielded by multifactor risk models, we developed a Lagrangian relaxation–based procedure that produces good quality solutions to portfolio optimization problems with a limit on the number of assets. The procedure also produces tight bounds on the optimal value of the objective function for such problems. Based on the procedure, a branch-and-bound algorithm was developed. This allows us to solve small- to medium-sized problems in reasonable time with branch-and-bound and, for larger problems, to produce good solutions accompanied with conservative estimates of their proximity to optimality. Computational results show the superiority of this method over both the diversity heuristic and branch-and-bound based on QP relaxation with binary variables methods.

ENDNOTES

The authors would like to thank Neil Gilfedder, Dan Stefek, and Miguel Alvarez for their help and advice in preparing this article for publication.

[1]For an even more detailed description, see Shaw, Liu, and Kopman [2008].

[2]The linear term c appearing in the objective function is equal to the asset return vector α in the absence of a benchmark. When the benchmark h_b is present, the objective expression $\alpha'h - \lambda(h - h_b)'(XFX' + D)(h - h_b)$ transforms into $(\alpha' + 2\lambda h_b'(XFX' + D))h - \lambda h'(XFX' + D)h$ (the constant term $\lambda h_b'(XFX' + D)h_b$ can be omitted), so in this case $c = \alpha' + 2\lambda h_b'(XFX' + D)$.

[3]For a description of a subgradient algorithm, see, for instance, Nemhauser and Wolsey [1999].

REFERENCES

Charpin, F., and D. Lacaze. "Using Binary Variables to Obtain Small Optimal Portfolios." *Journal of Portfolio Management,* Fall 2007, pp. 68–72.

Jansen, R., and R. Van Dijk. "Optimal Benchmark Tracking with Small Portfolios." *Journal of Portfolio Management,* Winter 2002, pp. 33–39.

Nemhauser, G., and L. Wolsey. *Integer and Combinatorial Optimization.* New York: John Wiley and Sons, 1999.

Rosenberg, B. "Extra-Market Components of Covariance in Securities Markets." *Journal of Financial and Quantitative Analysis,* 9 (1974), pp. 263–274.

Shaw, D., S. Liu, and L. Kopman. "Lagrangian Relaxation Procedure for Cardinality-Constrained Portfolio Optimization." *Optimization Methods and Software,* 23 (2008), pp. 411–420.