Portfolio construction through mixed integer programming

Dimitris Bertsimas[*], Christopher Darnell,[†] and Robert Soucy[‡]

Sloan WP #4024
May 1998

[*]Boeing Professor of Operations Research, Sloan School of Management and Operations Research Center, Rm. E53-363, Massachusetts Institute of Technology, Cambridge, Mass. 02139.
[†]Grantham, Mayo, van Otterloo & Co. LLC, 40 Rowes Wharf, Boston, Mass. 02110.
[‡]Grantham, Mayo, van Otterloo & Co. LLC, 40 Rowes Wharf, Boston, Mass. 02110.

Portfolio construction through mixed integer programming

Dimitris Bertsimas * Christopher Darnell† Robert Soucy ‡

May 1998

Abstract

We propose mixed integer programming (MIP) methods to construct a portfolio that is "close" (in terms of sector and security exposure) to a target portfolio, has the same liquidity, turnover and expected return with the target portfolio, controls frictional costs, and does so with fewer distinct stocks and requires fewer transactions. We also apply MIP methods to a portfolio consisting of several sub-portfolios. The algorithm has been implemented at Grantham, Mayo and van Otterloo & Co. LLC., (GMO) a leading investment management firm currently employing 170 people worldwide and managing over $26 billion. The MIP approach is currently being used in the construction of 11 quantitatively managed portfolios representing over $8 billion in assets since October, 1996. The benefits from the implementation of the project include: **(a)** Keeping the existing client business. As an example, GMO kept one $400 million separate account because the MIP algorithm reduced the number of positions from 1300 to 400-500 with careful control of turnover and transaction costs. **(b)** Making possible important new growth opportunities. Two U.S. small-market capitalization funds were successfully launched at year-end 1996 using the MIP portfolio construction process. **(c)** The number of names has been reduced by an average 40%-60% with only a marginal decrease in liquidity, while maintaining the turnover, performance, and sector exposures of the target portfolio. **(d)** The annual cost of trading the portfolios has decreased by at least $4 million due to a 75%-85% reduction in the number of trading tickets written to trade the portfolios. Additional savings of the same order of magnitude will be achieved when the MIP methods are applied to the large stock U.S. funds. **(e)** Significant improvement of the trading process was achieved. **(f)** A dramatic improvement in performance in simulation in a U.S. fund consisting of growth stocks with small market capitalization. If this result continues to hold for other funds, and we expect that it will, the use of the technology will have a dramatic impact on the future investment returns and success of the entire quantitative group at GMO.

*Dimitris Bertsimas, Boeing Professor of Operations Research, Sloan School of Management and Operations Research Center, Rm. E53-359, Massachusetts Institute of Technology, Cambridge, Mass 02139.

†Christopher Darnell, Grantham, Mayo, van Otterloo & Co. LLC, 40 Rowes Wharf, Boston, Mass. 02110.

‡Robert Soucy, Grantham, Mayo, van Otterloo & Co. LLC, 40 Rowes Wharf, Boston, Mass. 02110.

1 Introduction

After the seminal work of Markowitz (1959), portfolio managers routinely use quadratic programming methods to construct large scale portfolios. In the classical theory of portfolio optimization, portfolio managers select the fraction of wealth $w(i)$ invested in stock i in order to minimize the risk of the portfolio, measured by the variability of the return, which is a quadratic function of the decision variables, subject to linear constraints (one of the constraints is typically the requirement that the expected return of the portfolio is at least a certain target). This approach takes into account nicely the tradeoff between risk and return.

In the practice of portfolio construction, however, there are complications that the classical theory does not address. In particular, it is quite common for the number of different stocks (names) in the portfolio to be very large. Moreover, in the process of rebalancing the portfolio, the number of transactions (trading tickets) can also be large. The combination of a large number of names and a large number of tickets increases the costs of trading as both custodial fees and transaction costs increase. It is therefore desirable to construct portfolios that minimize the number of names in the portfolio as well as the number of tickets.

This paper summarizes a project that has been developed and implemented in the investment firm Grantham, Mayo and van Otterloo & Co LLC. (GMO) that uses mixed integer programming methods to construct a portfolio that is "close" (in terms of sector and security exposure) to a target portfolio, has the same liquidity, turnover and expected return as the target portfolio, controls frictional costs, and does so with fewer names and requires fewer tickets. Although the use of quadratic programming methods in the construction of portfolios is well documented in the academic literature and is widely used in practice, we are not aware of any use of mixed integer programming methods in the construction of portfolios in practice. Moreover, to the best of our knowledge, the problem we report here has not been addressed in the academic literature.

The rest of the paper is structured as follows. Section 2 gives background about the firm and the project. Section 3 describes the mixed integer programming approach that we used in the construction of portfolios. Section 4 describes the implementation of the algorithm. Section 5 reports simulation and actual results from the implementation of the project. Section 6 discusses the project's impact on the operations at GMO.

2 Background Information

Grantham, Mayo and van Otterloo & Co. LLC. (GMO), founded in 1977, is an investment management firm that has over $26 billion in assets under management and has over 170 employees worldwide. Its clients consist of pension funds, educational endowments, foundations, and a few large international organizations. The firm offers a wide range of mutual funds (both equity and fixed income funds) in the US and international markets.

Over the last decade, the most rapidly growing area of the firm has been the quantitative investment group. This group uses computer systems to design, implement and trade stock and bond portfolios for large institutional clients. The quantitative effort has begun in 1984 to provide investment services to clients in addition to that which could be provided by the traditional investment areas of the firm. The successful traditional investment groups, which depend on a small group of professionals to research investment ideas, quickly reached capacity limits for assets under management and have had little or no ability to accept new client funds since 1984. A quantitative investment process, with its ability to analyze thousands of securities using a variety of investment techniques and to rigorously control trading costs, can have a much larger capacity ceiling. The quantitative investment process at GMO makes extensive use of statistical, simulation, and optimization methodologies to meet return targets of clients with controlled risk in a large capacity format.

By 1996, the quantitative group had $15 billion in assets, with a sophisticated institutional client base that included leading investment institutions such as the World Bank, the International Monetary Fund, GTE, IBM, Harvard, Yale, and Princeton. The group had offices in Boston, San Francisco, London, and Sydney, managing money in every stock and bond market in the world. A record of above average and consistent investment returns had earned the confidence of a large number of institutions despite widespread discomfort with quantitative investment techniques in general. In 1996, GMO received the Most Innovative Award from Global Investor, an organization that ranks investment managers, for the most innovative investment manager of the year.

It was also clear in 1996, however, that a number of serious threats were developing to the continued success of the largest part of the quantitative group, the division managing stocks. The principle investment strategy of the quantitative stock division is a style known as value investing, which compares the prices of individual securities, of groups of securities like industries, and of entire country stock markets to theoretical values derived from economic and statistical models.

Value investing was performing below expectations in many markets around the world, and as a consequence the investment returns of the some of the largest and most important quantitative funds were trending down.

The response in the quantitative research group was to develop a technology that could improve the firm's capabilities in two critical areas: diversification of investment styles and control of the portfolio and trade construction process. An investment style represents a particular investment philosophy; for example a value investment style attempts to predict security returns using value related characteristics, like price to book or price to earnings ratios. In order to implement a particular investment style to forecast security returns, multivariate linear and non-linear regression models, called multifactor models, are used that have the characteristics of the particular investment style as independent variables. Diversification of investment style had been an ongoing effort at GMO, but now there was a need for a technology that could support a large-scale and comprehensive multiple investment style process. Multiple investment styles have been discussed in the financial literature and implemented at some other quantitative investment firms. While general multifactor models do provide diversification compared to a single factor model, there typically is no clear relationship between the performance of the individual factors in a multifactor model and the performance of a composite portfolio constructed using the model. In the worst case, for example, it is possible for the return of the composite portfolio to be lower than the return of the worst performing portfolio constructed using a single factor. This is certainly not the type of diversification that a client would find acceptable.

Large clients, in fact, typically divide their funds within a given asset class, like U.S. stocks, among at least several investment managers who have distinct investment styles. This insures that the composite investment result will be a linear combination of the underlying investment styles. In this framework the ex-ante and ex-post composite mean return is a simple weighted average of the ex-ante and ex-post composite mean returns of the individual managers. The ex-ante risk of the composite fund, as represented by the ex-ante variance of the composite return, is always less than the risk of the average manager.

Clients have widespread acceptance of this linear diversification framework, and the GMO quantitative group uses this framework to provide this style of diversification within individual funds. This is accomplished by partitioning portfolios into distinct "sub-portfolios," each with a distinct investment style. As an example, the U.S. Core Fund, a large companies U.S. stock fund, uses four sub-portfolios: a value, momentum, neglect and a cash flow sub-portfolio. As another example, in the International Core Fund, a large foreign countries fund, each country uses 2-5 investment styles that work best in that country, based on simulation and actual investment results. Each combination of a country and investment style gives fifty distinct sub-portfolios, each of which represents an average of only 2% of the total. The whole process leads to vast diversification as well as decreased transaction costs from the ability of sub-portfolios to cross with each other before trading in the market. This benefit is not present in the traditional investment framework, in which a client uses distinct external managers.

The GMO quantitative group follows the industry practice of constructing portfolios using quadratic optimization. Optimizing multiple sub-portfolios simultaneously is not a problem that has been addressed by others. This type of optimization presents the technical challenge of increasing the dimension of the problem by a factor roughly equal to the number of sub-portfolios, since, in general, each sub-portfolio can own any of the securities that the total portfolio can own. In the International Core portfolio, for example, there are fifty distinct sub-portfolios, which increases the dimension of the problem substantially.

The most difficult technical requirement we faced was a practical limit on the number of securities a composite portfolio could hold. The problem was most acute for GMO's large international equity portfolios. Increasing diversification into non-value based sub-portfolios was producing composite portfolios that had 1000-1500 different securities. Industry norms for comparable international portfolios are closer to 200 securities. This was producing problems in three critical areas:

1. **Client confidence was eroding.** The confluence of disappointing returns along with an apparently never ending increase in the number of securities in GMO's portfolios was leading clients to believe that the quantitative group was gradually losing control of the investment process. Several key clients were threatening to leave the firm.

2. **Operational complexity was straining the capacities of the trading, settlements, and accounting groups.** Portfolios with thousands of securities produced trade lists with many thousands of trading tickets to be executed, settled, and accounted for in markets around the world.

3. **There was widespread concern regarding high operational costs.** Ever increasing operational costs like ticket charges and custodial costs were of increasing concern to both

GMO top management and to many clients, since the costs are shared by both groups.

A simple way to reduce the number of names in a portfolio is to add a post-processing step to the quadratic portfolio optimization. A post-processing procedure that eliminates both security positions and security trades smaller than some threshold can reduce the number of positions and trade tickets to any desired level. There are two reasons why this approach was and continues to be infeasible at GMO.

The first is that small positions in a portfolio tend to be positions in securities with small market capitalization. Accompanying the underperformance of value investing was the underperformance of small market capitalization stocks around the world. Statistical research at GMO has demonstrated that small market capitalization stocks have underperformed. Their resulting undervaluation, has led GMO to the strategic forecast that small stocks will outperform over the next several years. A procedure that tends to eliminate small market capitalization stocks is clearly unacceptable in this environment.

The other serious drawback of this post-processing approach is a problem in all market environments: decreased control over the portfolio construction and trade creation process. Post-processing that eliminates hundreds of positions will interfere substantially with optimization objectives and constraints on key variables like expected returns, risk, sub-portfolio allocations, and transaction costs. Although clients might not care about globally optimal versus multiple-step portfolio construction, they do demand complete control and accountability from GMO.

In addition to these serious issues for existing portfolios, in 1996 GMO wanted to prepare to take advantage of the growing opportunity in small market capitalization stocks by developing and simulating a portfolio construction process for a series of small market capitalization stock funds. Preliminary results for even a single country (U.S.) small market capitalization stock fund showed that GMO's existing multiple-sub-portfolio process was going to produce composite portfolios with more than 1500 securities. In light of the above discussion, this was not going to be acceptable. It was clear that GMO had to develop a global optimization method that could jointly optimize multiple sub-portfolios and control the number of positions and trades in the composite portfolio. Since the number of positions and trades in the composite portfolio is an inherently integer quantity, we decided that we would use mixed integer programming methods to globally optimizing the portfolio.

3 The mixed integer programming approach

In this section, we describe the methodology we followed first for a single portfolio and then for a portfolio that involves several sub-portfolios.

3.1 The single portfolio problem

We consider in this section a single portfolio, which is rebalanced monthly. Let $\mathbf{w}_0 = (w_0(1), w_0(2), \ldots, w_0(N))$ be the current portfolio, where $w_0(i)$ is the fraction of the portfolio invested in stock $i = 1, \ldots, N$. Let $\mathbf{w}_t = (w_t(1), w_t(2), \ldots, w_t(N))$ be the target portfolio, i.e., the portfolio that it is desirable to own after rebalancing. In some funds, in which we applied the mixed integer programming methods, this target portfolio is constructed using quadratic optimization techniques. In some other funds the weight $w_t(i)$ for stock i in the target portfolio is a closed form expression of the predicted return $\alpha(i)$ and the market capitalization of stock i.

The objective is to decide the final portfolio $\mathbf{w}_f = (w_f(1), w_f(2), \ldots, w_f(N))$ that satisfies the following characteristics:

1. **Closeness of the final and the target portfolios.**
   Given that the target portfolio $\mathbf{w}_t$ was selected taking into account several factors (tradeoffs in risk, return and liquidity), it is desirable for the portfolio $\mathbf{w}_f$ to be as close as possible to $\mathbf{w}_t$. This requirement is captured by a term

   $$\sum_{i=1}^{N} |w_f(i) - w_t(i)|$$

   in the objective function.

2. **Sector exposure.**
   It is desirable that the exposure to different sectors in the economy between the target and final portfolios are as close as possible. Examples of sectors include utilities, financial firms, etc. This requirement is captured by a term

   $$\sum_{s=1}^{K} \lambda_{\text{sec}}(s) \left| \sum_{i=1}^{N} M_s(i)(w_f(i) - w_t(i)) \right|,$$

in the objective, where $M_s(i)$ is a 0/1 index denoting membership of stock i in sector s, and $\lambda_{\text{sec}}(s)$ is a user specified penalty for sector s that captures the importance of the requirement that the difference in sector exposure between the final and the target portfolios is small.

3. **Number of names.**

   It is desirable to have a portfolio with a small number of names. For this reason we define

$$y_{\text{names}}(i) = \begin{cases} 1, & \text{if } w_f(i) > 0, \\ 0, & \text{if } w_f(i) = 0. \end{cases}$$

   This requirement is captured by a term

$$\lambda_{\text{names}} \sum_{i=1}^{N} y_{\text{names}}(i)$$

   in the objective, where λ_{names} is a user specified penalty that captures the importance of the requirement that the number of names in the portfolio is small.

4. **Number of tickets.**

   It is desirable to have a portfolio with a small number of transactions. For this reason we define

$$y_{\text{tickets}}(i) = \begin{cases} 1, & \text{if } |w_f(i) - w_0(i)| > 0, \\ 0, & \text{if } |w_f(i) - w_0(i)| = 0. \end{cases}$$

   Note that a transaction is made, and therefore a trading ticket is written, if there is a difference between the current and the final portfolio. This requirement is captured by a term

$$\lambda_{\text{tickets}} \sum_{i=1}^{N} y_{\text{tickets}}(i)$$

   in the objective, where λ_{tickets} is a user specified penalty that captures the importance of the requirement that the number of tickets written is small.

5. **Return of the portfolio.**

It is desirable to have a final portfolio with high return. This is captured by a term

$$-\lambda_\alpha \sum_{i=1}^{N} \alpha(i) w_f(i),$$

where $\alpha(i)$ is the expected return of stock i, and λ_α is a user specified penalty that captures the importance of the requirement that the return of the portfolio is high. The reason of the negative sign is that the overall objective is minimization.

6. **Liquidity of the portfolio.**

It is desirable to have a portfolio with high liquidity or equivalently with low illiquidity. In particular, as the position in a stock increases, it is harder, and thus more expensive, to trade it. Thus, the illiquidity of stock i is captured by a piecewise linear convex function $f_{\text{lq}}(i, w_f(i))$ depicted in Figure 1.

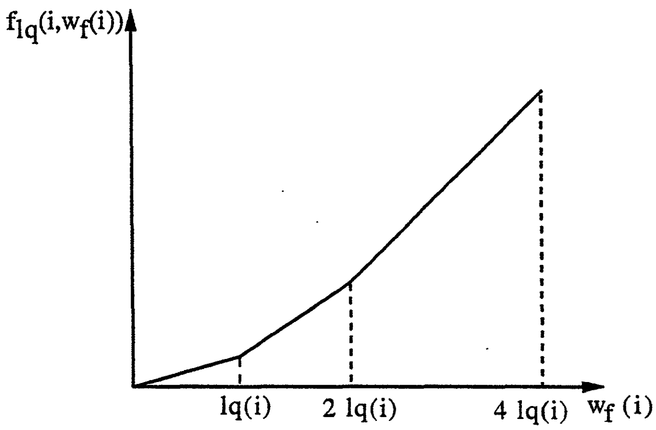

Figure 1: The illiquidity function.

For every stock i there is an illiquidity index $\text{lq}(i)$, such that the illiquidity function is modeled as

$$f_{\text{lq}}(i,x) = \begin{cases} \text{ls}(1,i)x, & 0 \le x \le \text{lq}(i), \\ \text{ls}(2,i)(x - \text{lq}(i)) + \text{ls}(1,i)\text{lq}(i), & \text{lq}(i) \le x \le 2 \cdot \text{lq}(i), \\ \text{ls}(3,i)(x - 2 \cdot \text{lq}(i)) + \text{ls}(2,i)2 \cdot \text{lq}(i), & 2 \cdot \text{lq}(i) \le x \le 4 \cdot \text{lq}(i). \end{cases}$$

The liquidity consideration is captured by a term

$$\lambda_{\text{illiquidity}} \sum_{i=1}^{N} f_{\text{lq}}(i, w_f(i)),$$

where $\lambda_{\text{illiquidity}}$ is a user specified penalty that captures the importance of the requirement that the illiquidity of the portfolio is low.

7. **Transaction costs.**

   It is desirable to minimize total transaction costs. In particular, as the trading position increases relative to the daily volume of a stock, transaction costs are higher as there is an impact in the market. Clearly, as the traded amount increases, the price impact is greater, and thus the effect increases. The transaction cost from trading stock i is captured by a piecewise linear convex function $f_{\text{tc}}(i, |w_f(i) - w_0(i)|)$ depicted in Figure 2.

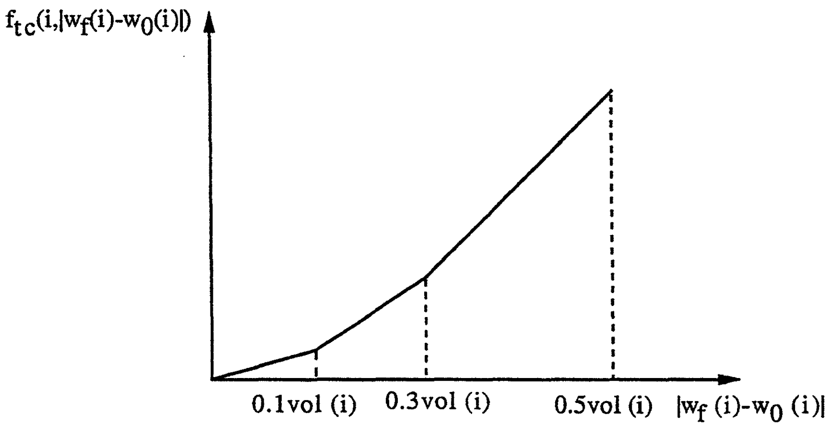

Figure 2: The transaction cost function.

For every stock i there is a volume index $\text{vol}(i)$, such that the transaction cost function is modeled as

$$f_{\text{tc}}(i,x) = \begin{cases} \text{cs}(1,i)x, & 0 \le x \le 0.1 \cdot \text{vol}(i), \\ \text{cs}(2,i)(x - 0.1 \cdot \text{vol}(i)) + \text{cs}(1,i)0.1 \cdot \text{vol}(i), & 0.1 \cdot \text{vol}(i) \le x \le 0.3 \cdot \text{vol}(i), \\ \text{cs}(3,i)(x - 0.3 \cdot \text{vol}(i)) + \text{cs}(3,i)0.3 \cdot \text{vol}(i), & 0.3 \cdot \text{vol}(i) \le x \le 0.5 \cdot \text{vol}(i). \end{cases}$$

The transaction cost consideration is captured by a term

$$\lambda_{\text{tc}} \sum_{i=1}^{N} f_{\text{tc}}(i, |w_f(i) - w_0(i)|),$$

where λ_{tc} is a user specified penalty that captures the importance of the requirement that the transaction cost is small.

The complete formulation is presented in Appendix A. The model uses various penalties, denoted by λ, that capture the relative importance of the various objectives. The penalties are chosen heuristically after extensive experimentation. We run the algorithm for a given set of penalties, and observe the performance of the portfolio in historical simulations over approximately twenty years, re-optimizing monthly. If a characteristic of the portfolio is not considered satisfactory its corresponding penalty is increased. Thousands of runs are performed in order to determine satisfactory penalties.

3.2 The multiple portfolio problem

We have also applied the mixed integer programming methods to a portfolio that consists of several sub-portfolios indexed by j, $j = 1, \ldots, S$. Let $w_0(i,j)$, $w_t(i,j)$ be the current and target position respectively of stock i in sub-portfolio j, $i = 1, \ldots, N$, $j = 1, \ldots, S$. The objective is to decide the final position $w_f(i,j)$ of stock i in sub-portfolio j. Clearly, the current, target, and final position of stock i in the portfolio is

$$w_0(i) = \sum_{j=1}^{S} w_0(i,j), \quad , w_t(i) = \sum_{j=1}^{S} w_t(i,j), \quad w_f(i) = \sum_{j=1}^{S} w_f(i,j).$$

As before the objective is to decide the final weights $w_f(i,j)$ so that the final and target portfolios are close, the total sector exposure is similar, the number of names in each sub-portfolio, as well as the total portfolio, is small, the total number of trading tickets written is small, the total return of the portfolio and its liquidity are high, and the total transaction costs are small.

Using similar methodology, as in the single portfolio problem, we formulated a mixed integer programming model. One of the significant advantages of constructing a portfolio that consists of several sub-portfolios that represent different investment philosophies is that when the portfolio is optimized, the transaction costs can be significantly reduced as the sub-portfolios trade among each other without the fund incurring transaction costs. This is one of the major attractions for clients who in this way have the benefit of diversification among different investment philosophies, while having global control of transaction costs. This point is further amplified in Section 5.1, where a historical simulation during 1982-1997 of a U.S. small growth optimized portfolio consisting of three sub-portfolios demonstrates that the optimized portfolio outperforms significantly the target portfolio.

4 Implementation

The mixed integer programming model was implemented in FORTRAN using CPLEX 4.0 as the underlying mixed integer programming solver. The model runs on a Digital Equipment Corp. (DEC) Alpha cluster running OPENVMS and Dell PentiumPRO PC systems running Microsoft WindowsNT 4.0 with an x-windows, telnet, nfs, and ftp connection to the local office Alpha cluster. The Alpha cluster component used for optimization is an AlphaServer 4100 5/466 including four 466 MHz CPUs (acting independently, not in parallel) and two GB of main memory. External storage is well over 150 GB.

As we have already mentioned, the software optimization engine is CPLEX 4.0, a product available through the CPLEX division of ILOG. In particular, CPLEX provides a callable library of routines. The callable library is at the heart of an application we call TRGTOPT (target optimization) which consists of approximately 10,000 lines of Fortran90 code written at GMO. TRGTOPT is command line driven accepting as input a text file that describes the specifics of the problem to be solved. Once the input file is parsed into a dynamic data structure, TRGTOPT queries the database for target portfolio and sub-portfolio information and sets up the problem to be solved by CPLEX.

Improvements to the implementation and problem formulation have been considerable since the initial implementation in October 1996. At that time, when we were trying to solve the single portfolio problem (without any sub-portfolios) with approximately 1500 securities (the number of variables is typically 8-10 times the number of securities), it could take as long as 15 hours of CPU time to solve a **single problem**. This was considered unsatisfactory, as GMO's desire was to solve many problems for simulation purposes. Typically, before introducing a new method in the portfolio construction process, the quantitative group attempts to simulate this new method historically. During the fall and winter months of 1996, we tried to decrease these running times substantially so that a simulation over twenty years that rebalances the portfolio monthly (240 problems need to be solved) was feasible.

After extensive experimentation with parameters associated with CPLEX like node and variable selection strategies, setting branching priorities, and adjusting the stopping criteria we were able to improve the solution times considerably. Further improvements were realized by strengthening the formulation, and thus improving the linear programming relaxation bounds. In summary, the most important factors that contributed in the improvement of the running times are:

1. **Strengthening the formulation to improve the relaxation bounds.**
   We describe some of the improvements in Appendix B.

2. **Using the structure of the problem to set node selection and branching priorities.**
   Given the target vector, it is unlikely that the largest positions in the target portfolio will be eliminated, while it is more likely that the smallest positions will be eliminated. So, we first select to branch on the variables that correspond to the largest positions in the target portfolio. Regarding branching priorities, for the largest positions we branch first on the option to keep the position, while for the smallest positions, we branch first on the option to eliminate the position.

3. **Experimentation with the various parameters of CPLEX.**
   The stopping criterion, and the sub-optimality allowed had an effect on the running time.

We now routinely solve a 1500 security problem in a few minutes which enables us to set up simulations that run overnight solving hundreds of problems. Table 1 shows the sizes of four problems and the CPU times needed to solve each problem.

5 Results

In this section, we first describe simulation results that were performed prior to the implementation of the mixed integer programming approach, and were used to convince top management at GMO of the effectiveness of the method. We then present results of actual implementations in eleven portfolios of total market value $ 8.158 billion.

Portfolio	Average size rows/columns	max size rows/columns	problems	CPU/problem (minutes)
SMALG	1084/8884	NA	1	1.4
SMALV	1592/11153	NA	1	1.8
UK	598/4982	864/7411	241	0.4
small growth	5434/41969	6546/51532	181	4.0

Table 1: Examples of portfolios run using the mixed integer programming model. The results for the first two portfolios SMALG (small growth) and SMALV (small value) are from an actual run of the mixed integer programming model. The results for the UK portfolio and the small growth portfolio represent historical simulations over 241 and 181 months respectively. A summary of the results for these simulation experiments are shown in Figures 3 and 5 respectively.

5.1 Simulation Results

Prior to the implementation of the mixed integer programming approach extensive historical simulations were performed. In this section, we present two examples of such simulations that were performed recently that illustrate that our methods can solve truly large problems, and can lead to significant improvements in performance. These two examples also show that the mixed integer programming algorithm can be used in either "tracking mode" or "performance enhancing mode."

Historical simulation of the UK portfolio

The first example is of a United Kingdom (UK) portfolio that is the weighted sum of three sub-portfolios. The three sub-portfolios reflect independent models for adding value to the UK market. In this example, the penalties were chosen so that the optimized portfolio tracks the performance of the target portfolio as well as the performance of the target sub-portfolios but does it with many fewer securities. We refer to this mode of running the algorithm as "tracking mode." Referring to Figures 3, 4, the relative strength shows the performance of the target portfolio and the optimized portfolio relative to the benchmark (in this case, the MSCIP UK index). The scale reflects a cumulative benchmark multiple, so that a value of 1.7 (the approximate value in 1997) means that cumulative portfolio performance is 1.7 times that of the benchmark. Since the benchmark returned 2000% (in round numbers) cumulatively (16.4% annually) over the 20 year period, the implication is that the portfolio returned 3400% cumulatively (19.5% annually). We can observe that the optimized portfolio tracks the target portfolio very closely. The average number of securities in the optimized portfolio is approximately 55% (80/146) that of the target portfolio, while the average number of tickets was reduced by 60% in the optimized portfolio. Note that the average monthly turnover is virtually identical (the turnover spikes are due to a sub-portfolio that is traded only once per year). Figure 4 shows that the optimized three sub-portfolios track very closely the target sub-portfolios as well.

Historical simulation of the small growth US portfolio

The second example is of a US portfolio consisting of growth stocks with small market capitalization. It shows that the optimized portfolio actually **outperforms** the target portfolio **significantly**. Figure 5 shows that the optimized portfolio over the period 1982-1997 has cumulative performance 2 times that of the benchmark, while the target portfolio has 1.6 times that of the benchmark. Note that the other characteristics are similar to the previous example: a 55% reduction in the number of names, an 80% reduction in the number of tickets, and virtually identical turnover. We have achieved this significant improvement in performance by adjusting the various penalties λ that are present in the objective function. We refer to this mode of running the algorithm as "performance enhancing mode." While these results on performance have only been seen in simulation, they have been instrumental in convincing top management in GMO (a) to apply this methodology extensively throughout the firm and (b) to launch new funds using mixed integer programming methods.

5.2 Implemented Results

The mixed integer programming methodology has been applied to eleven portfolios of total market value \$ 8.158 billion. Table 2 shows the names of these portfolios, the corresponding market value, the number of names before and after running the mixed integer programming model, and the number of tickets before and after running the model. The mixed integer programming method has been applied to these eleven portfolios in the period October, 1996-January, 1997. Table 2 shows an average reduction of 48.7% in the number of names and an average reduction of 79.3% in the number of trading tickets.

Portfolio Name	Mkt Value ($ billion)	# Names Before	# Names After	% Names Reduction	# Tickets Before	# Tickets After	% Tickets Reduction
ISF	4308	1285	772	39.9	764	186	75.7
CHIC	651	1172	699	40.4	613	134	78.1
GTEI	369	1061	626	41.0	552	123	77.7
IMFQ	387	1003	606	39.6	576	127	78.0
IQNT	475	1066	467	56.2	652	140	78.5
SCAP	251	893	516	42.2	571	137	76.0
AMER	288	1188	417	64.9	553	102	81.6
JSF	212	373	222	40.5	231	49	78.8
SMALG	386	986	394	60.0	917	156	83.0
SMALV	743	1451	579	60.1	1298	206	84.1
QVF	88	866	424	51.0	671	129	80.8

Table 2: Quantitative portfolios constructed using the mixed integer programming approach.

6 Impact

The impact on GMO's quantitative group of the mixed integer programming development has been significant in many areas. In general, this technology has made a significant contribution both to keeping GMO's existing client business and to making possible important new growth opportunities. In GMO's international portfolios, the number of positions has decreased by 40-65%, and the number of trades requiring processing decreased by 75-85%. GMO's clients understand this has been accomplished without compromising GMO's multiple investment goals. In one $400 million international separate account, reducing the number of positions from 1300 to 400-500 was a condition for keeping the account. Our mixed integer programming process achieved this goal, and accomplished the reduction with careful control of turnover and transaction costs according to a timetable specified by the client.

Two U.S. small-market capitalization funds were successfully launched at year-end 1996 using the mixed integer programming portfolio construction process. The investment returns of the funds and the control over the portfolio and trading process have met or exceeded our expectations. A year after launch the funds have a value of $1.1 billion, a clear success for the quantitative group and the mixed integer programming technology.

Reductions in operational and trading costs have also met our expectations. Annual savings of about $4 million have been realized from the sharp drop in trading tickets in the international and small stock U.S. funds. Additional savings of the same order of magnitude will be achieved when the mixed integer methods are applied to the large stock U.S. funds. Although there are no variable cost savings from the large reduction in number of positions in the international and small stock U.S. funds, GMO's fixed price custodial contracts have avoided potentially large increases in annual rates.

Internal efficiencies are reflected in the firm's not having to add new trading operations personnel in the last two years, in the decline in trade instruction and settlement errors in the international area, and in a marked reduction in stress in trading and trading-related areas within the firm. The ability to optimize the large international trades has given us the ability to better control frictional transaction costs. These are the costs incurred by large investors as security prices "move away" from buy and sell orders. When these costs are poorly controlled and therefore large, they can have a serious negative impact on total investment returns. While difficult to quantify, the international trading group has reported a significant decrease in the difficulty and cost of executing international trades.

Since its development in 1996, we have used the mixed integer technology in "tracking mode," with the goal of tracking as closely as possible target portfolios and sub-portfolios subject to control over number of positions, number of trades, and transaction costs. Substantial improvements in processing times for the mixed integer programming algorithm have enabled the exploration of goals more complex than simple tracking. One very significant result has come from simulations that relax the tracking objective to improve the ex ante portfolio return. As discussed in Section 5, this approach has been examined in detail for the U.S. small stock growth fund simulation, and has significantly enhanced simulated returns of the fund without violating other constraints.

This "free lunch" comes from the remarkable ability of the global optimization of multiple sub-portfolios to sharply increase turnover at the sub-portfolio level without increasing total portfolio turnover. With sub-portfolio turnover increased from 90% per year to 200% per year, sub-portfolio returns are significantly enhanced. Total portfolio returns are guaranteed to be enhanced proportionately because of GMO's linear framework, and with the global optimization effectively arranging crossing trades between the sub-portfolios, total portfolio turnover is held to target levels. If this result continues to hold for funds other than the small stock growth fund, and we expect that it will, this use of the technology will have a dramatic impact on the future investment returns and success of the entire quantitative group at GMO.

Appendix A: A Complete Formulation of a Single Portfolio Problem

The complete formulation of the single portfolio problem is presented below. In addition to the variables $w_f(i)$, $y_{\text{names}}(i)$, $y_{\text{tickets}}(i)$ already defined in Section 3, the formulation uses auxiliary variables $y(i)$, $f(i)$, $x(s)$, to model $|w_f(i) - w_t(i)|$, $|w_f(i) - w_0(i)|$, and $\left|\sum_{i=1}^{N} M_s(i)(w_f(i) - w_t(i))\right|$ respectively. It also uses the auxiliary variables $x_1(i)$, $x_2(i)$, $x_3(i)$, to model the different pieces of the piecewise linear, and convex illiquidity function, and the auxiliary variables $z_1(i)$, $z_2(i)$, $z_3(i)$ to model the different pieces of the piecewise linear, and convex transaction cost function.

$$
\begin{aligned}
\text{minimize} \quad & \sum_{i=1}^{N} y(i) + \sum_{s=1}^{K} \lambda_{\text{sec}}(s) x(s) + \lambda_{\text{names}} \sum_{i=1}^{N} y_{\text{names}}(i) \\
& + \lambda_{\text{tickets}} \sum_{i=1}^{N} y_{\text{tickets}}(i) - \lambda_{\alpha} \sum_{i=1}^{N} \alpha(i) w_f(i) \\
& + \lambda_{\text{illiquidity}} \sum_{i=1}^{N} \Big(ls(1,i)x_1(i) + ls(2,i)x_2(i) + ls(3,i)x_3(i) \Big) \\
& + \lambda_{\text{tc}} \sum_{i=1}^{N} \Big(cs(1,i)z_1(i) + cs(2,i)z_2(i) + cs(3,i)z_3(i) \Big) \\
\text{subject to} \quad & \sum_{i=1}^{N} w_f(i) = 1, \\
& w_f(i) - w_t(i) \leq y(i), && \forall i \\
& -(w_f(i) - w_t(i)) \leq y(i), && \forall i \\
& w_f(i) - w_0(i) \leq f(i), && \forall i \\
& -(w_f(i) - w_0(i)) \leq f(i), && \forall i \\
& x(s) \geq \sum_{i=1}^{N} M_s(i)(w_f(i) - w_t(i)), && \forall s \\
& x(s) \geq -\sum_{i=1}^{N} M_s(i)(w_f(i) - w_t(i)), && \forall s \\
& w_f(i) \leq y_{\text{names}}(i), && \forall i \\
& f(i) \leq y_{\text{tickets}}(i), && \forall i \\
& w_f(i) = x_1(i) + x_2(i) + x_3(i), && \forall i \\
& f(i) = z_1(i) + z_2(i) + z_3(i), && \forall i
\end{aligned}
$$

$$
\begin{array}{ll}
0 \le x_1(i) \le \mathrm{lq}(i), & \forall i \\
0 \le x_2(i) \le \mathrm{lq}(i), & \forall i \\
0 \le x_3(i) \le 2 \cdot \mathrm{lq}(i), & \forall i \\
0 \le z_1(i) \le 0.1 \cdot \mathrm{vol}(i), & \forall i \\
0 \le z_2(i) \le 0.2 \cdot \mathrm{vol}(i), & \forall i \\
0 \le z_3(i) \le 0.2 \cdot \mathrm{vol}(i), & \forall i \\
y(i), w_f(i), f(i) \ge 0, & \forall i \\
x(s) \ge 0, & \forall s \\
y_{\mathrm{names}}(i), y_{\mathrm{tickets}}(i) \in \{0,1\}, & \forall i.
\end{array}
$$

Appendix B: Formulation Enhancements

The constraint $w_f(i) \le y_{\mathrm{names}}(i)$ has been used in the formulation presented in Appendix A to relate the continuous and the discrete variables. In practice, however, no final weight is expected to be larger than a threshold, say 5%. Moreover, as the market capitalization of a stock decreases, this threshold will be smaller. For this reason, we strengthen the constraint $w_f(i) \le y_{\mathrm{names}}(i)$ by replacing it by the constraint:

$$w_f(i) \le a_i \, w_{\max} \, y_{\mathrm{names}}(i),$$

where $w_{\max}$ is the maximum weight of a stock in the portfolio, and a_i are constants that depend on the market capitalization of the stock. The resulting relaxation is stronger than the one outlined in Appendix A.

Acknowledgements

We would like to thank Professor Stephen Graves, editor of the special issue, and Professor Richard Rosenthal for insightful comments that improved the paper.

References

[1] Markowitz, H., *Portfolio Selection: Efficient Diversification and Investments*, John Wiley, New York, 1959.

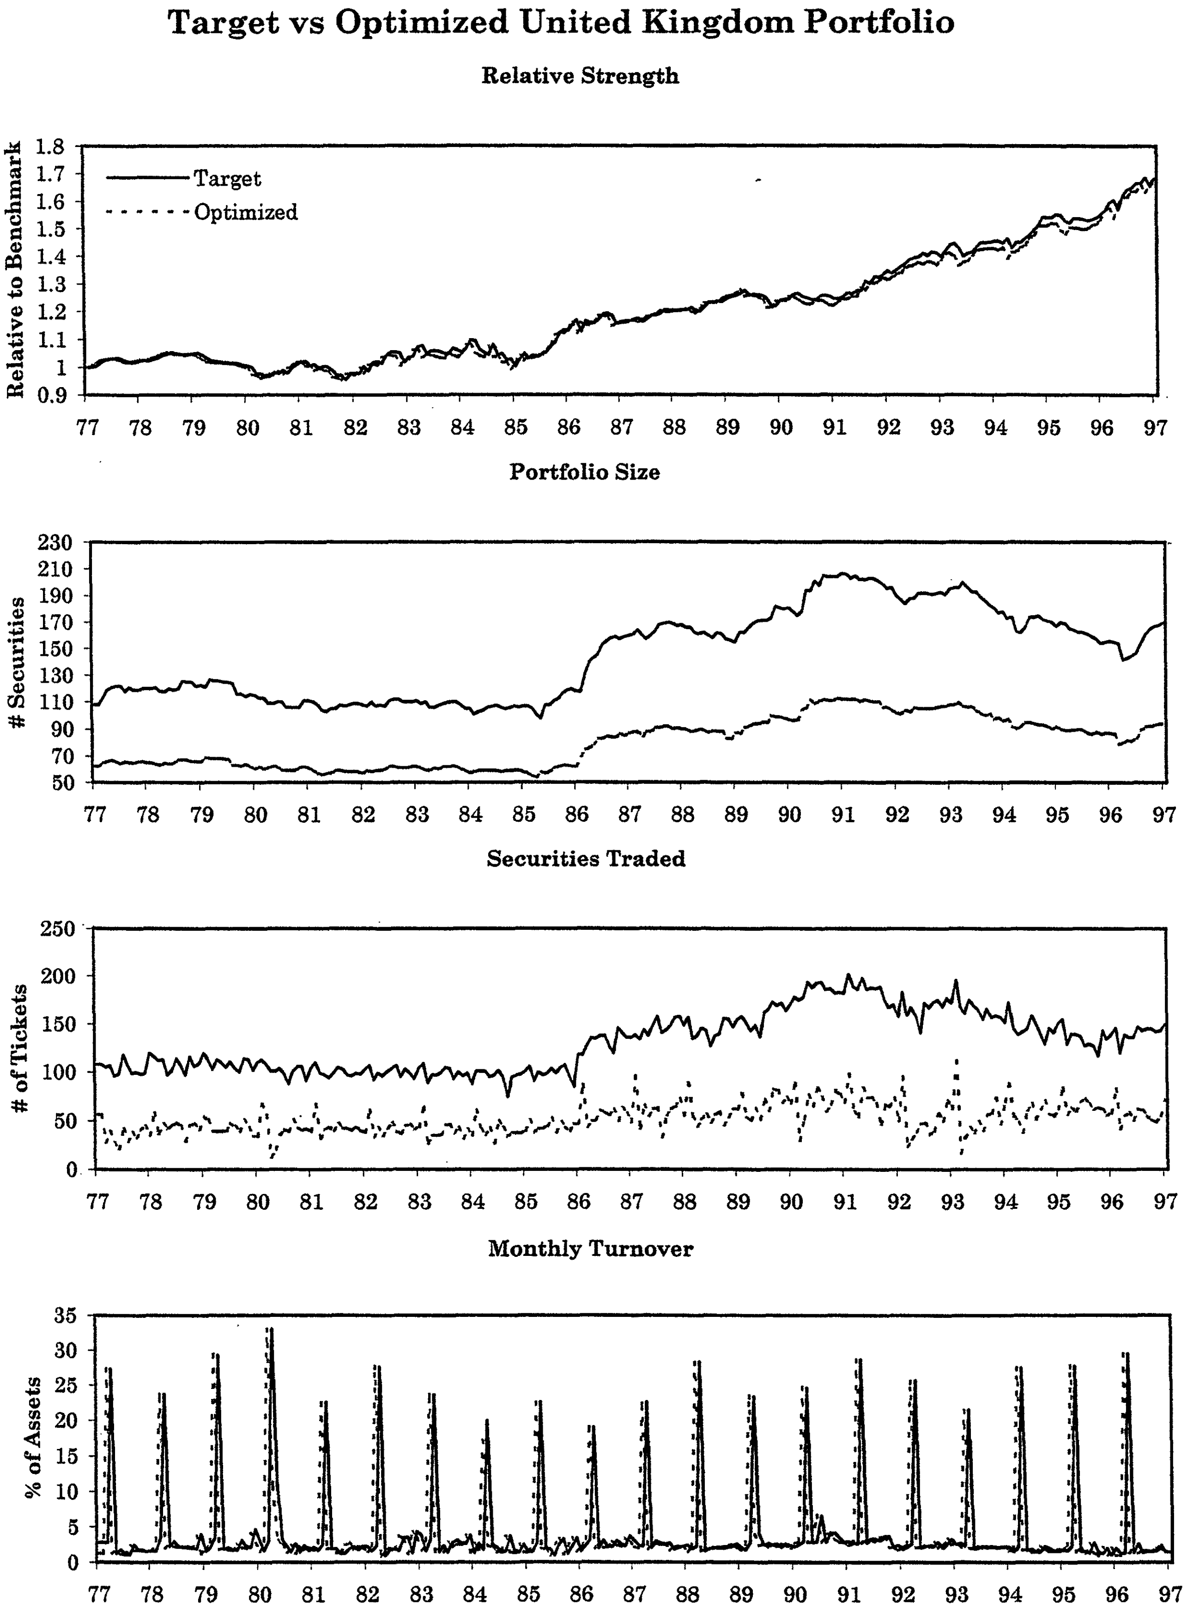

Figure 3: Simulated performance of the UK portfolio (1977-1997)

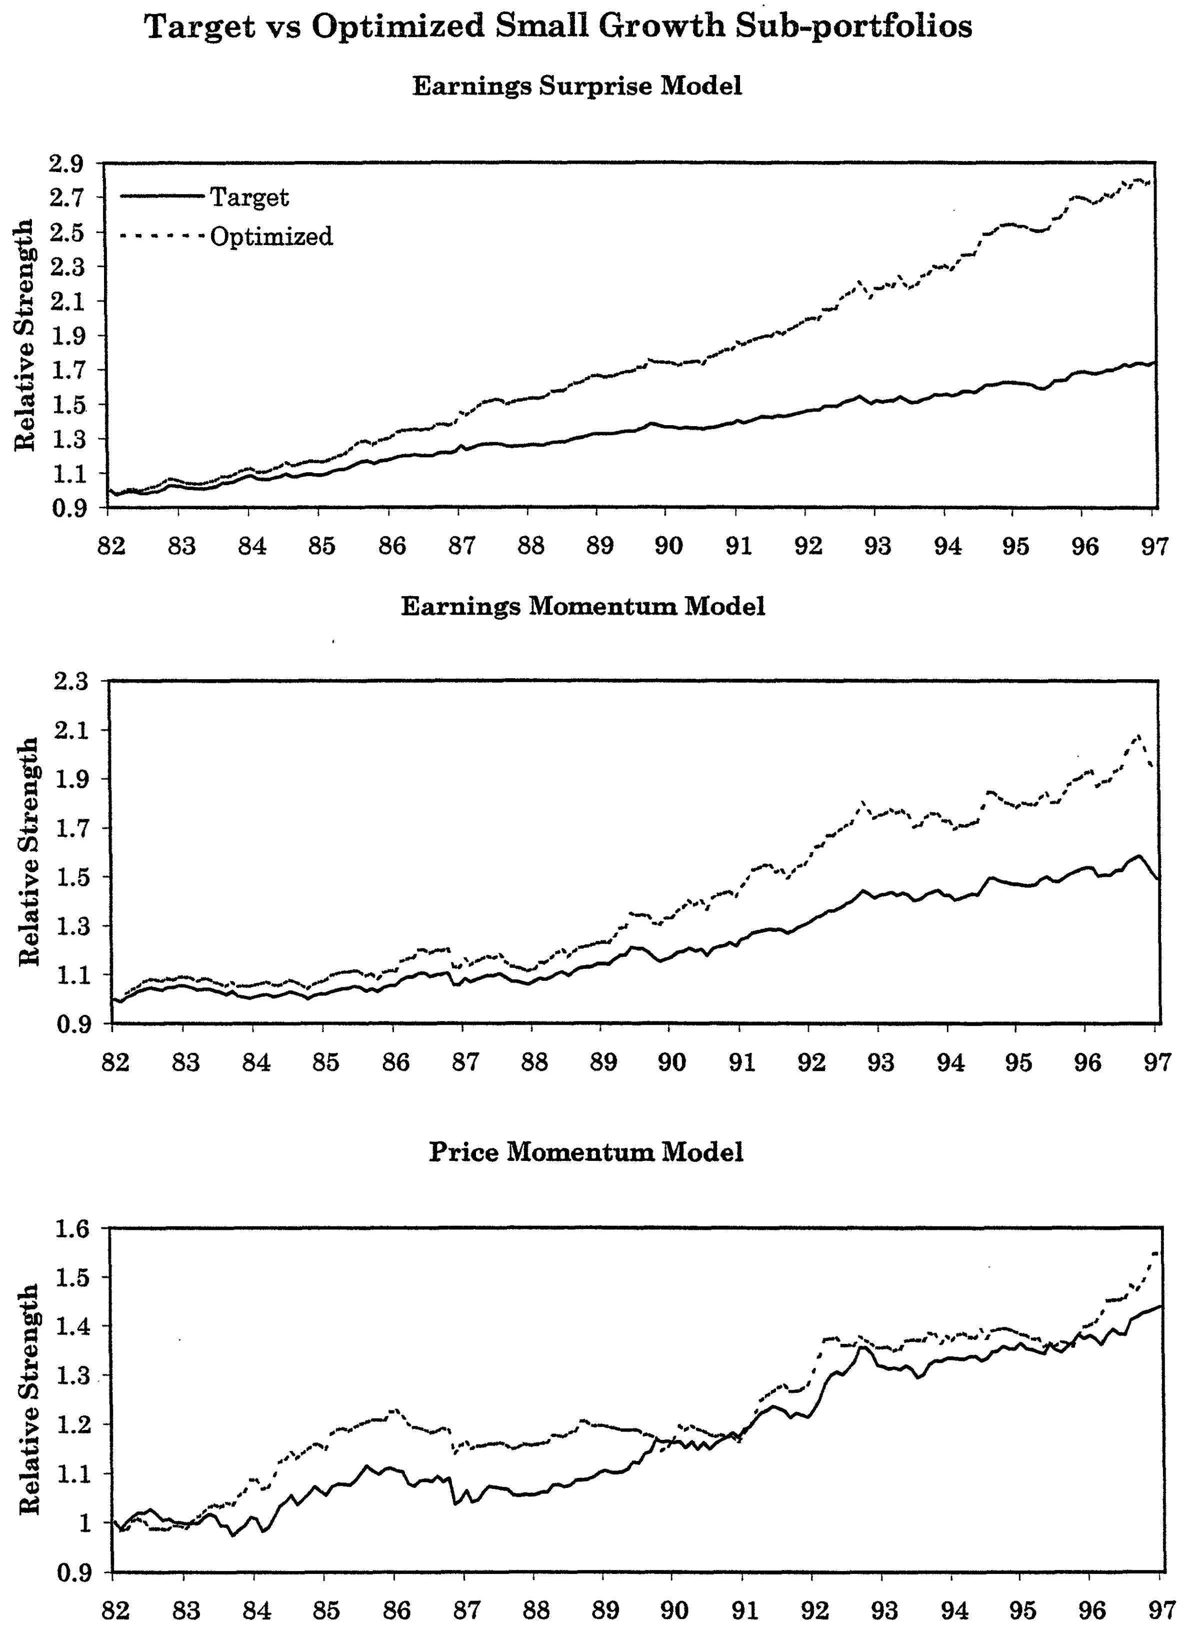

Figure 6: Simulated performance of the three small growth sub-portfolios (1982-1997)