

Portfolio Construction Through Mixed-Integer Programming at Grantham, Mayo, Van Otterloo and Company

Dimitris Bertsimas

*Sloan School of Management and Operations
Research Center
Massachusetts Institute of Technology,
Room E53-363
Cambridge, Massachusetts 02142*

Christopher Darnell

*Grantham, Mayo, Van Otterloo and Company LLC
40 Rowes Wharf
Boston, Massachusetts 02110*

Robert Soucy

Grantham, Mayo, Van Otterloo and Company LLC

Grantham, Mayo, Van Otterloo and Company LLC (GMO) uses mixed-integer-programming (MIP) methods to construct portfolios that are close (in terms of sector and security exposure) to target portfolios, have the same liquidity, turnover, and expected return as the target portfolios, control frictional costs, and do so with fewer distinct stocks and with fewer transactions. It also applies MIP methods to portfolios consisting of several subportfolios. It uses the MIP approach to construct 11 quantitatively managed portfolios representing over \$8 billion in assets. The benefits from this implementation include (1) keeping the existing client business; (2) making possible important new growth opportunities; (3) reducing the number of stock names by an average 40 to 60 percent; (4) reducing the annual cost of trading the portfolios by at least \$4 million by reducing the number of trading tickets written by 75 to 85 percent; (5) improving the trading process; and (6) improving performance in simulation in a US fund consisting of growth stocks with small market capitalization.

Since the seminal work of Markowitz [1959], portfolio managers have routinely used 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 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 trade-off between risk and return.

GMO has over \$26 billion in assets under management.

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 because both custodial fees and transaction costs increase. It is therefore desirable to construct portfolios that minimize the number of names in the portfolio and the number of tickets.

We developed and implemented a method in the investment firm Grantham, Mayo, Van Otterloo and Company LLC (GMO) that uses mixed-integer-programming methods to construct a port-

folio 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 fewer tickets. Although the use of quadratic-programming methods in the construction of portfolios is well documented in the academic literature and they are 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 addressed has not been described previously in the academic literature.

Background Information

Grantham, Mayo, Van Otterloo and Company 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 firm began its quantitative effort in 1984 to provide investment services to clients in addition to those that the traditional investment areas of the firm could provide. 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.

Increasing diversification into nonvalue-based subportfolios was producing composite portfolios that contained 1,000 to 1,500 different securities.

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 innova-

tive 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 principal 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, such as 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 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 process of constructing portfolios and trades. An investment style represents a particular investment philosophy; for example, a value investment style attempts to predict security returns using value-related characteristics, such as price-to-book or price-to-earnings ratios. To implement a particular investment style to forecast security returns, firms use multivariate linear and nonlinear regression models, called multifactor models, that have the characteristics of the particular investment style as independent variables. GMO has made an ongoing effort to diversify investment style, but it now needed 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 single-factor models, 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, the return of the composite portfolio can 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 (US stocks, for example) among at least several investment managers who have distinct investment styles. This ensures 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 generally accept this linear diversification framework, and the GMO quantitative group uses it to provide this style of diversification within individual funds. It does this by partitioning portfolios into distinct subportfolios, each with a distinct investment style. As an example, the US Core Fund, a stock fund consisting of large US companies, uses four subport-

folios: a value, a momentum, a neglect, and a cash-flow subportfolio. As another example, in the International Core Fund, a stock fund consisting of large international companies, GMO uses two to five investment styles in each country that work best in that country, based on simulation and actual investment results. The combinations of a country and an investment style make up 50 distinct subportfolios, each of which represents an average of only two percent of the total. The whole process leads to vast diversification as well as to decreased transaction costs because the subportfolios 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 subportfolios 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 subportfolios, since, in general, each subportfolio can own any of the securities that the total portfolio can own. In the International Core portfolio, for example, there are 50 distinct subportfolios, 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

nonvalue-based subportfolios was producing composite portfolios that contained 1,000 to 1,500 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) Concern regarding high operational costs was widespread. Ever-increasing operational costs, such as ticket charges and custodial costs, were of increasing concern to both GMO top managers and to many clients, since both shared the costs.

The technology will have a dramatic impact on the investment returns of the entire quantitative group.

A simple way to reduce the number of names in a portfolio is to add a postprocessing step to the quadratic optimization of the portfolio. A postprocessing 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. For two reasons, this approach was and continues to be infeasible at GMO: First, 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 large stocks 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 postprocessing approach is a problem in all market environments: decreased control over the portfolio-construction and trade-creation process. Postprocessing that eliminates hundreds of positions will interfere substantially with optimization objectives and constraints on such key variables as expected returns, risk, subportfolio 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 (US) small-market-capitalization stock fund showed that GMO's existing multiple-subportfolio process was going to produce composite portfolios with more than 1,500 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 subportfolios 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 optimize the portfolio.

The Mixed-Integer-Programming Approach

We first consider a single portfolio, which is rebalanced monthly. Let $\mathbf{w}_0 = (w_0(1), w_0(2), \dots, w_0(N))$ be the current portfolio, where $w_0(i)$ is the fraction of the portfolio invested in stock $i = 1, \dots, N$. Let $\mathbf{w}_t = (w_t(1), w_t(2), \dots, w_t(N))$ be the target portfolio, that is, 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 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 $\mathbf{w}_f = (w_f(1), w_f(2), \dots, w_f(N))$ that has the following characteristics: (1) closeness to the target portfolio, (2) exposure to different economic sectors close to that of the target portfolio, (3) a small number of

names, (4) a small number of transactions, (5) high return, (6) high liquidity, and (7) low transaction costs.

(1) Given that the target portfolio \mathbf{w}_t was selected taking into account several factors (trade-offs 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) It is desirable that the exposure to different sectors in the economy between the target and final portfolios be as close as possible. Examples of sectors include utilities and financial firms. 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 be small.

(3) 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 be small.

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

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

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 be small.

(5) 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 be high. The reason for the negative sign is that the overall objective is minimization.

(6) 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 $f_{\text{liq}}(i, w_f(i))$ depicted in Figure 1.




Figure 1: The illiquidity function.

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

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

The liquidity consideration is captured by a term

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

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

(7) 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 become higher because portfolio trades have 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.

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




Figure 2: The transaction cost function.

function is modeled as

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

The transaction-cost consideration is captured by a term

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

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

The complete formulation is presented in the appendix. The model uses various penalties, denoted by λ , that capture the relative importance of the various objectives. We choose penalties 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 20 years, reoptimizing monthly. If we consider a characteristic of the portfolio unsatisfactory, we increase its corresponding penalty. We perform thousands of runs to determine satisfactory penalties.

The Multiple-Portfolio Problem

We have also applied mixed-integer-programming methods to a portfolio that

consists of several subportfolios indexed by $j, j = 1, \dots, S$. Let $w_0(i, j), w_t(i, j)$ be the current and target position, respectively, of stock i in subportfolio $j, i = 1, \dots, N, j = 1, \dots, S$. The objective is to decide the final position $w_f(i, j)$ of stock i in subportfolio j . Clearly, the current, target, and final position of stock i in the portfolio is

$$w_0(i) = \sum_{j=1}^S w_0(i, j),$$

$$w_t(i) = \sum_{j=1}^S w_t(i, j),$$

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

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 subportfolio, 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 methodology similar to that for the single-portfolio problem, we formulated a mixed-integer-programming model. One of the major advantages of constructing a portfolio that consists of several subportfolios that represent different investment philosophies is that, in optimizing the portfolio, one can reduce the transaction costs by having the subportfolios trade with each other and thus avoid incurring transaction costs for the fund. This is a major attraction for clients. They have the benefit of diversification among different investment philosophies, while having global control of transaction costs.

Portfolio	Average size rows/columns	Maximum size rows/columns	Problems	CPU/problem (minutes)
SMALG	1,084/8,884	NA	1	1.4
SMALV	1,592/11,153	NA	1	1.8
UK	598/4,982	864/7,411	241	0.4
Small growth	5,434/41,969	6,546/51,532	181	4.0

Table 1: The table shows the sizes of four problems and the CPU times needed to solve each problem. In these four 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 (Figure 3) and the small growth portfolio (Figure 5) represent historical simulations over 241 and 181 months, respectively.

Implementation

We implemented the mixed-integer-programming model 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 Windows NT 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 Alpha Server 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.

The software optimization engine CPLEX 4.0 is 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 FORTRAN 90 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.

We have made many improvements to the implementation and problem formulation since the initial implementation in October 1996. At that time, when we were trying to solve the single-portfolio problem (without any subportfolios) with approximately 1,500 securities (the number of variables is typically equal to eight to 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 20 years that rebalances the portfolio monthly (240 problems to be solved) was feasible (see Table 1).

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. We realized further improvements by strengthening the formulation and thus improving the linear-programming relaxation bounds. The most important factors that contributed to the improvement of the running times are the following:

- (1) We strengthened the formulation to improve the relaxation bounds (appendix).
- (2) We used 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) We experimented with the various parameters of CPLEX and found that the stopping criterion and the suboptimality allowed had an effect on the running time.

We now routinely solve a 1,500-security problem in a few minutes, which enables us to set up simulations that run overnight solving hundreds of problems.

Results

Prior to implementing the mixed-integer-programming approach, we performed extensive historical simulations. We used these to convince top management at GMO of the effectiveness of the method. We performed two such simula-

tions recently that show 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 subportfolios. The three subportfolios reflect independent models for adding value to the UK market. In this example, we chose the penalties so that the optimized portfolio tracks the performance of the target portfolio as well as the performance of the target subportfolios but does it with many fewer securities. We refer to this mode of running the algorithm as tracking mode. Figures 3 and 4 show 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 2,000 percent (in round numbers) cumulatively (16.4 percent annually) over the 20-year period, the implication is that the portfolio returned 3,400 percent cumulatively (19.5 percent annually). The optimized portfolio tracks the target portfolio very closely. The average number of securities in the optimized portfolio is approximately 55 percent (80/146) that of the target portfolio, while the average number of tickets was reduced by 60 percent in the optimized portfolio. The average monthly turnover is virtually



Figure 3: The figure shows an example of running the method in the tracking mode. It compares the simulated performance of the target and optimized United Kingdom portfolio during the period 1977–1997. The portfolio consists of three subportfolios. The first panel shows the ratio of the return of each of the two portfolios to the return of the benchmark. The second and third panels show the number of securities and number of trading tickets in both the target and the optimized portfolios. The fourth panel shows the turnover for both portfolios.

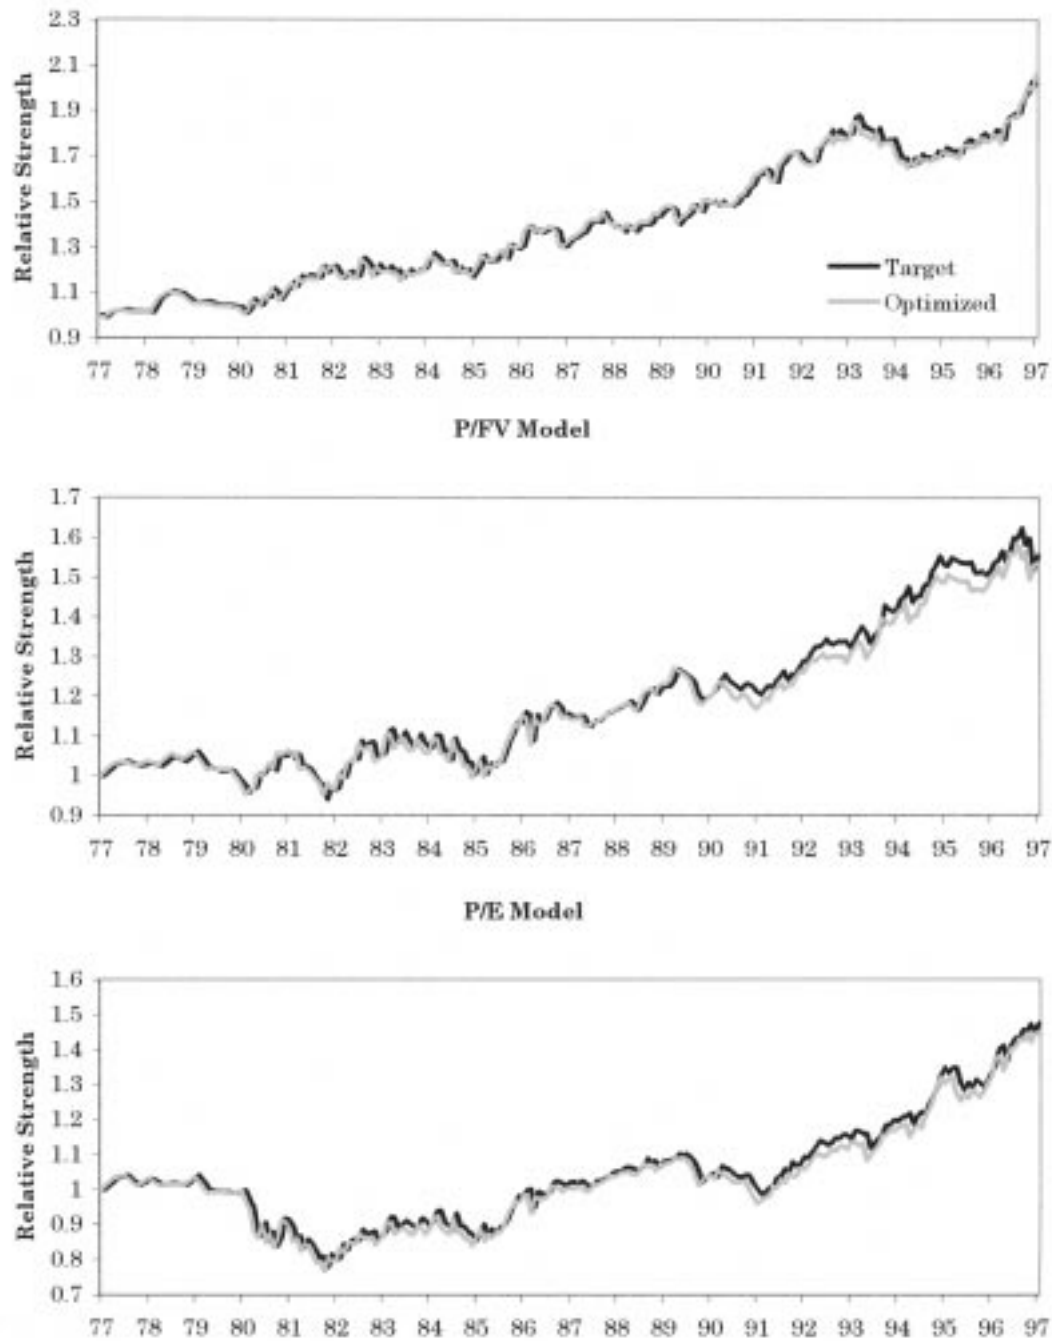


Figure 4: The figure compares the simulated performance of the target and optimized subportfolios for the UK portfolio during the period 1977–1997. The three subportfolios use a momentum, a price to fair value, and a price to earnings model, respectively.

identical (the turnover spikes are due to a subportfolio that is traded only once per year). Figure 4 shows that the optimized three subportfolios track the target subportfolios very closely 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 (Figures 5 and 6). It shows that the optimized portfolio actually outperforms the target portfolio significantly. The optimized portfolio over the period 1982 through 1997 has cumulative performance twice that of the benchmark, while the target portfolio has a performance 1.6 times that of the benchmark (Figure 5). The other characteristics are similar to the previous example: a 55-percent reduction in the number of names, an 80-percent reduction in the number of tickets, and virtually identical turnover. We have achieved this improvement in performance by adjusting the various penalties present in the objective function. We refer to this mode of running the algorithm as performance-enhancing mode. While we have seen these results only in simulation, they have been instrumental in convincing GMO top management (a) to apply this methodology extensively throughout the firm and (b) to launch new funds using mixed-integer-programming methods.

Implemented Results

We have applied the mixed-integer-programming methodology to portfolios with a total market value of \$ 8.158 billion in the period October 1996 to January 1997 (Table 2). We achieved an average reduction of 48.7 percent in the number of

names and an average reduction of 79.3 percent in the number of trading tickets.

Impact

The mixed-integer-programming development has had a major impact on GMO's quantitative group in many areas. In general, this technology has contributed both to keeping GMO's existing client business and to making possible new growth opportunities. In GMO's international portfolios, the number of positions has decreased by 40 to 65 percent, and the number of trades requiring processing has decreased by 75 to 85 percent. GMO's clients understand that GMO has accomplished this without compromising its investment goals. In one \$400 million international separate account, reducing the number of positions from 1,300 to 400 to 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 the client specified.

GMO launched two US small-market-capitalization funds successfully at the end of 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. GMO has realized annual savings of about \$4 million from the sharp drop in trading tickets in the international and small stock

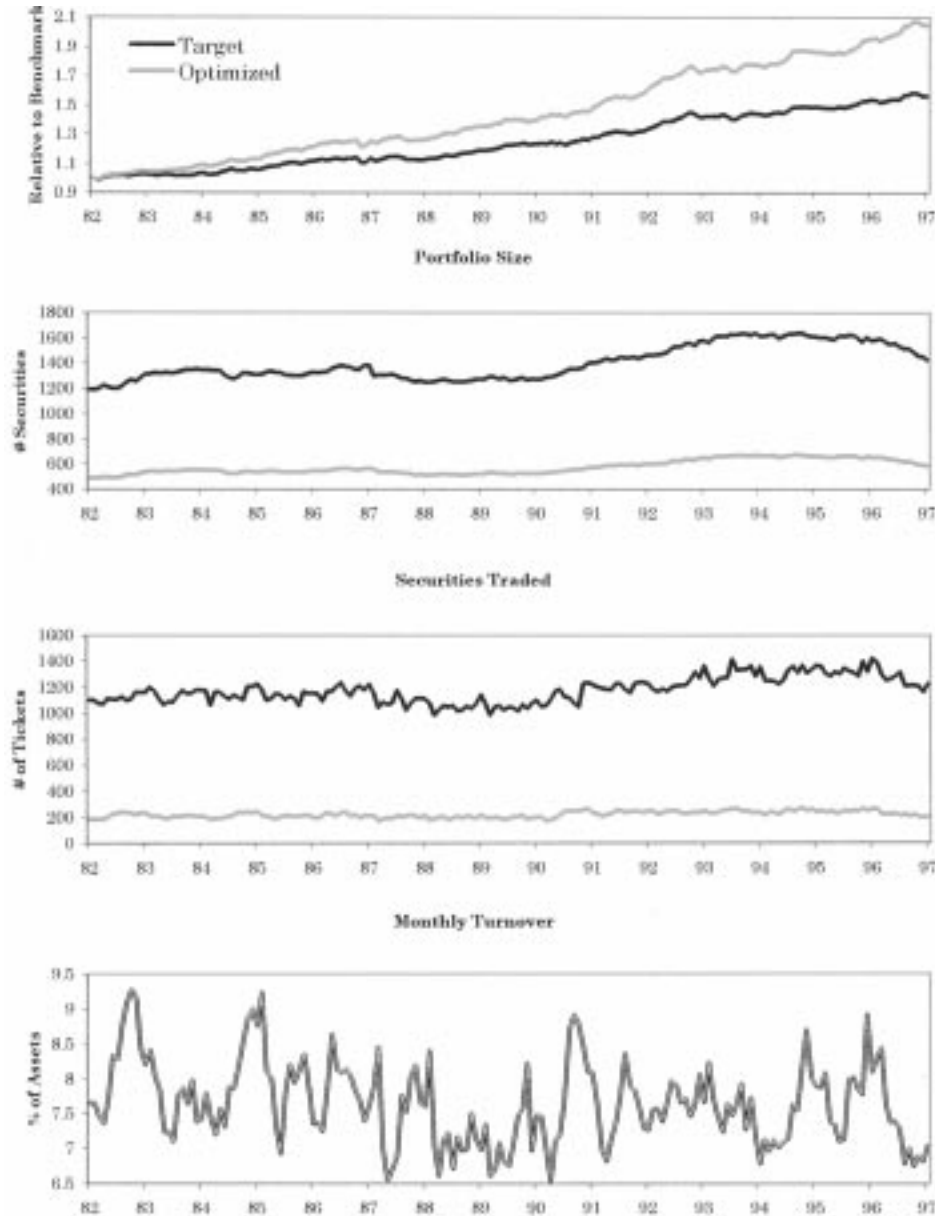

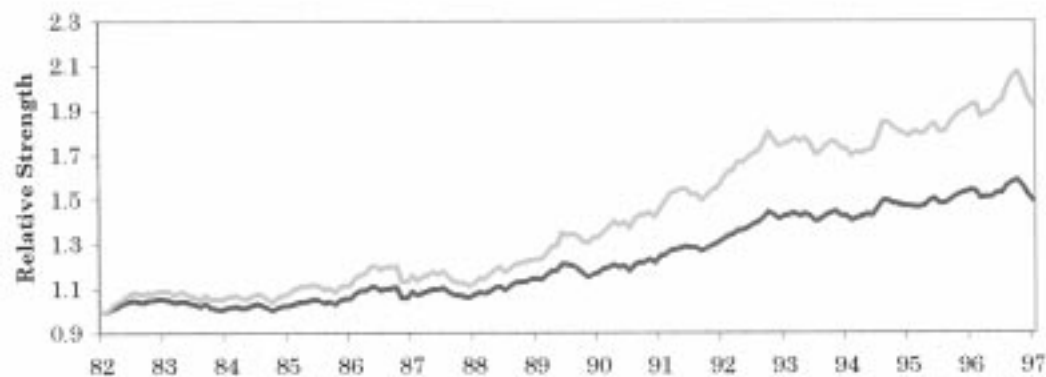


Figure 5: The figure shows an example of running the method in the performance-enhancing mode. It compares the simulated performance of the target and optimized US small-growth portfolio during the period 1982–1997. The portfolio consists of three subportfolios. The first panel shows the ratio of the return of each of the two portfolios to the return of the benchmark. The second and third panels show the number of securities and number of trading tickets in both the target and the optimized portfolios. The fourth panel shows the turnover for both portfolios.



Earnings Momentum Model



Price Momentum Model




Figure 6: The figure compares the simulated performance of the target and optimized subportfolios for the US small-growth portfolio during the period 1982-1997. The three subportfolios use an earnings surprise, earnings momentum, and price momentum model, respectively.

US funds. It will achieve additional savings of the same order of magnitude when it applies the mixed-integer-programming methods to the large stock US funds. Although GMO has achieved no savings in variable costs by reducing the number of positions in the international and small stock US funds, it has avoided potentially large increases in annual rates by charging fixed prices for its custodial contracts.

Because of the internal efficiencies achieved, the firm has not had to add new trading-operations personnel in the last two years, trade instruction and settlement errors in the international area have declined, and stress in trading and trading-related areas within the firm shows a marked reduction. 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 the prices specified in the buy and sell orders. When these costs are poorly controlled and therefore large, they

can drastically decrease total investment returns. The international trading group has reported a decrease in the difficulty and cost of executing international trades, a decrease that is difficult to quantify.

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 subportfolios subject to control over number of positions, number of trades, and transaction costs. Improvements in processing times for the mixed-integer-programming algorithm have enabled us to explore goals more complex than simple tracking. One important result has come from simulations that relax the tracking objective to improve the ex-ante portfolio return. We have examined this approach in detail for simulations of the US small stock growth fund and find that it enhances simulated returns of the fund without violating other constraints.

This free lunch comes from the remarkable ability of the global optimization of

Portfolio name	Market value (\$ billion)	No. names before	No. names after	% names reduction	No. tickets before	No. tickets after	% tickets reduction
ISF	4,308	1,285	772	39.9	764	186	75.7
CHIC	651	1,172	699	40.4	613	134	78.1
GTEI	369	1,061	626	41.0	552	123	77.7
IMFQ	387	1,003	606	39.6	576	127	78.0
IQNT	475	1,066	467	56.2	652	140	78.5
SCAP	251	893	516	42.2	571	137	76.0
AMER	288	1,188	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	1,451	579	60.1	1298	206	84.1
QVF	88	866	424	51.0	671	129	80.8

Table 2: The quantitative portfolios constructed using the mixed-integer-programming approach show improvements in market value and reductions in the number of names and the number of tickets.

multiple subportfolios to sharply increase turnover at the subportfolio level without increasing total portfolio turnover. Increasing subportfolio turnover from 90 percent per year to 200 percent per year greatly enhances subportfolio returns. GMO's linear framework guarantees that portfolio returns will be enhanced proportionately to the weight of each subportfolio because of GMO's linear framework, and with the global optimization effectively arranging crossing trades between the subportfolios, 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.

Acknowledgments

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

APPENDIX

A Complete Formulation of a Single-Portfolio Problem

The complete formulation of the single-portfolio problem is presented below. In addition to the variables $w_f(i)$, $y_{names}(i)$, $y_{tickets}(i)$ 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 $|\sum_{wi=1}^N M_s(i)(w_f(i) - w_t(i))|$, 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.

$$\begin{aligned}
 &\text{Minimize} \\
 &\sum_{i=1}^N y(i) + \sum_{i=1}^K \lambda_{sec}(s)x(s) \\
 &\quad + \lambda_{names} \sum_{i=1}^N y_{names}(i) \\
 &\quad + \lambda_{tickets} \sum_{i=1}^N y_{tickets}(i) - \lambda_\alpha \sum_{i=1}^N \alpha(i)w_f(i) \\
 &\quad + \lambda_{illiquidity} \sum_{i=1}^N (ls(1, i)x_1(i) + ls(2, i)x_2(i) \\
 &\quad + ls(3, i)x_3(i)) \\
 &\quad + \lambda_{tc} \sum_{i=1}^N (cs(1, i)z_1(i) + cs(2, i)z_2(i) \\
 &\quad + cs(3, i)z_3(i)) \\
 &\text{subject to} \\
 &\sum_{i=1}^N w_f(i) = 1, \\
 &w_f(i) - w_t(i) \leq y(i), \quad \forall i \\
 &-(w_f(i) - w_t(i)) \leq y(i), \quad \forall i \\
 &w_f(i) - w_0(i) \leq f(i), \quad \forall i \\
 &-(w_f(i) - w_0(i)) \leq f(i), \quad \forall i \\
 &x(s) \geq \sum_{i=1}^N M_s(i)(w_f(i) - w_t(i)), \quad \forall s \\
 &x(s) \geq -\sum_{i=1}^N M_s(i)(w_f(i) - w_t(i)), \quad \forall s \\
 &w_f(i) \leq y_{names}(i), \quad \forall i \\
 &f(i) \leq y_{tickets}(i), \quad \forall i \\
 &w_f(i) = x_1(i) + x_2(i) + x_3(i), \quad \forall i \\
 &f(i) = z_1(i) + z_2(i) + z_3(i), \quad \forall i \\
 &0 \leq x_1(i) \leq lq(i), \quad \forall i \\
 &0 \leq x_2(i) \leq lq(i), \quad \forall i \\
 &0 \leq x_3(i) \leq 2 \cdot lq(i), \quad \forall i \\
 &0 \leq z_1(i) \leq 0.1 \cdot vol(i), \quad \forall i \\
 &0 \leq z_2(i) \leq 0.2 \cdot vol(i), \quad \forall i \\
 &0 \leq z_3(i) \leq 0.2 \cdot vol(i), \quad \forall i \\
 &y(i), w_f(i), f(i) \geq 0, \quad \forall i \\
 &x(s) \geq 0, \quad \forall s \\
 &y_{names}(i), y_{tickets}(i) \in \{0, 1\}, \quad \forall i.
 \end{aligned}$$

Formulation Enhancements

The constraint $w_f(i) \leq y_{names}(i)$ has been used in the formulation to relate the continuous and the discrete variables. In practice, however, no final weight is expected to be larger than a threshold, say five percent. Moreover, as the market capitaliza-

tion of a stock decreases, this threshold will be smaller. For this reason, we strengthen the constraint $w_f(i) \leq y_{names}(i)$ by replacing it by the constraint:

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

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 here.

Reference

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

Jeremy Grantham, Founding Member, Grantham, Mayo, Van Otterloo and Company LLC, 40 Rowes Wharf, Boston, Massachusetts 02110, writes: "Grantham, Mayo, Van Otterloo and Company LLC (GMO) is an investment management firm founded in 1977. GMO currently employs 170 people worldwide and manages over \$26 billion for 653 clients. We have been taking an integer programming approach to the construction of 11 quantitatively managed portfolio representing over \$8 billion in assets since October 1996.

"Our portfolio construction methodology, employing multiple investment styles with diversification over time, lends itself to creating portfolios (called 'target' portfolios) that have desirable characteristics (low turnover, high liquidity, good performance) but contain a large number of securities. One way to reduce the number of securities would be to simply eliminate all of the small positions. That's typically not a bad first approximation but small positions are usually associated with smaller capitalization securities which we periodically

believe will outperform the larger capitalization securities and therefore want to maintain our small capitalization exposure, as we do today. Another motivation for an integer programming solution is to reduce the costs associated with trading the portfolio. We could account for frictional and brokerage costs in a quadratic programming problem but we also need to be able to control the number of tickets we have to write to trade the portfolio.

"So the optimization problem, and our motivation for pursuing integer programming methods, is to construct a portfolio that is 'close' (in terms of sector and security exposure) to the target portfolio, has about the same liquidity, turnover, and expected return, controls frictional costs, and does so with fewer names and requires fewer tickets.

"The benefits of employing an integer optimization framework are significant. First, the number of securities has been reduced by an average of 40–60 percent with only a marginal decrease in liquidity while maintaining the turnover, performance, and sector exposures of the target portfolio. Second, the costs of trading the portfolio have decreased by at least \$4 million due to a 75–85 percent reduction in the number of tickets written to trade it and the ability to have a cost model in the objective function."