Multi-stage Optimization for Long-term Investors

John M. Mulvey
Professor
Department of Operations Research and Financial Engineering
Bendheim Center for Finance
Princeton University
Princeton, NJ 08544

Mulvey@princeton.edu
http://www.princeton.edu/~bcf

September 2000

Abstract

Multi-stage simulation and optimization models are effective for solving long-term financial planning problems. Prominent examples include: asset-liability management for pension plans, integrated risk management for insurance companies, and long-term planning for individuals. Several applications will be briefly mentioned.

A multi-stage framework provides advantages over single-period myopic approaches. First, the investor gains an understanding of the risks that a long-term goal will be unfulfilled, such as retiring with adequate wealth. A multi-stage model can be more realistic than a single period model. Thus, assets such as equity, which reduce long-term risks while increasing short-term volatility, can be evaluated in a temporal setting. The tradeoff between long and short-term gains becomes apparent in a multi-period context. As a second advantage, enhanced returns are possible with dynamic investment strategies. For instance, the traditional approach of rebalancing assets to a fixed strategic benchmark generates higher returns when assets possess increased volatility. This “volatility pumping” is dampened by transaction and market impact costs. Only by solving a multi-stage optimization model can we discover the optimal rebalancing rules. Likewise, moving a large portfolio to a new strategic benchmark can be optimized. As a third example, individuals often hold assets with large embedded gains. Selling these assets triggers a capital gains tax. Again, these decisions can be evaluated by means of a multi-stage model. A real-world example from pension planning illustrates the concepts.

Three distinct approaches are available for solving the multi-stage optimization model: 1) dynamic stochastic control, 2) stochastic programming, and 3) optimizing a stochastic simulation model. We briefly review the pros and cons of these approaches; it seems unlikely that a single approach will dominate the others. We conclude with some topics for future research.

1. Introduction

The worldwide deregulation of markets gives rise to new investment opportunities such as hedge funds, private equity and venture capital, weather derivatives and catastrophic options on earthquakes/hurricanes. Unfortunately, many of these instruments possess greater volatility than traditional asset categories such as stocks and bonds and are difficult to hedge. Plus, the added volatility comes in bursts with severe consequences. Figure 1 depicts a distribution of losses for earthquakes in California. This long-tailed, asymmetric distribution is difficult to model by means of classical portfolio approaches.

Increasingly, investors are turning to multi-stage financial simulation and optimization models for assistance. These planning systems help the investor evaluate the consequences of decisions regarding asset allocation, business policy, capital structure, and other strategic choices. Importantly, a multi-stage model includes significant details that must be omitted from single-stage investment models.

Long-term investors can particularly benefit by a systematic multi-stage analysis. For example, individuals must choose an asset mix in order to generate retirement funds, or to pay for their children's college tuition. Most people have great difficulty understanding the impact of their investment decisions on the probability of attaining their retirement goals, for instance. In a similar fashion, pension plans must cope with achieving sufficient asset grow in order to meet the pension trust's funding obligations. Actuaries conduct annual reviews in order to measure funding surplus or deficit. The resulting asset and liability study provides information about the trust's ability to carry out its fiduciary duties. See Berger and Mulvey (1998), Cariño et al. (1994), Dempster (1998), Mulvey et al. (2000), and their references for further details. In these applications, multi-stage financial simulation and optimization models play an important role.

1.1 Market Impact Costs

A feature of newer financial markets, shared with other securities, is the tendency to display greater market impact costs[1] during severe movements and crises[2]. Often, we can attribute the degree of impact to the underlying economic environment. The 1998 summer/fall liquidity problem, for example, depicts the increase in market impact costs when investors attempt to cover their short positions on liquidity. Selling liquidity during a crisis can be hazardous to one's economic health.

Market impact costs arise in numerous guises. Taxes must be paid when assets are sold, or when dividends and interest are issued; individuals go to great lengths to avoid these impact costs. Other costs are less tangible but no less real. Selling a large stake in a small company is a second example. In certain cases, the investor must maintain their stake in an activity. Investors may wish to pull back on venture capital, for example, after receiving $60\text{-}70^{+}$ % annual gains over the past few years. However, it may be difficult to reduce exposure in venture capital investments earlier than the original plans. Here again, market impact costs may prevent an investor from re-balancing his portfolio in a timely fashion.

---

[1] We define all transaction costs by the term "market impact" costs.
[2] Of course, changing impact costs are nothing new. Anyone attempting to sell real estate during a major downturn understands the problem.

What are standard approaches for including market impact costs in financial planning models? A simple rule of thumb is to limit investments in illiquid asset categories to a modest exposure, sticking to actively traded assets such as the SP500 and U.S. government bonds. Another approach is to place a limit on turnover per period. For example, allowing no more than 3% turnover at any single juncture. See Davis and Norman (1990) and Taksar et al. (1998) for theoretical results for addressing transaction costs. It is challenging to analyze these strategies outside the scope of a multi-stage financial model.

1.2 Asset Categories Possessing Fat Tails

Including asset categories with extreme losses (or gains) within a portfolio model requires a careful consideration of rare events. For instance, the catastrophic (Cat) insurance shown in Figure 1 will generate a loss in excess of $250 million every 200 years on average. The expected annual gain for the asset is equal to $8 million, and the asset will generate a profit 65 % of the time. In addition, the returns are generally independent of other asset returns; hence these are excellent diversifiers[3]. How much of these securities, if any, should be included in the portfolio of a long-term investor?

A single-period model will have difficulty with this question since it supplies little information regarding the investor's chances of reaching his long-term goals or fulfilling fiduciary obligations. We have found that a better indicator of risk involves the probability of the investor reaching his or her goals. We call the concept "Goals-at-Risk (GaR[tm]). For example, individuals can readily understand the idea that they have a 90% probability of retiring at 65 years old with income equal to 70% of their pre-retirement salary. This concept is a reference point for evaluating alternative investment and savings strategies.

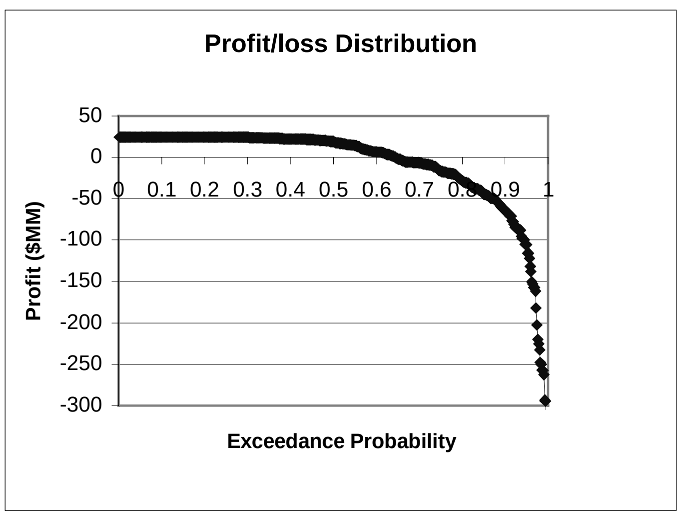

Figure 1
A Fat-Tailed Profit/loss Distribution for Earthquake Insurance

[3] Hedge funds and other non-traditional assets possess similar benefits.

In certain cases, optimization methods can be employed to reshape the profit/loss probability distribution, thereby improving the risk-adjusted returns. For instance, we have worked with insurance companies in the area of issuing policies for earthquakes (in California) and hurricane insurance (along the US East coast). As mentioned, this business possesses long-tailed loss distributions. By optimizing on risk-adjusted profit, we eliminate or reduce policies that are concentrating risks. Figure 2 shows an example. By reducing the business by less than 10 %, we improve the risk-adjusted profit from 17.6 % to 51.8% under one capital allocation rule, and from 13.6 % to 32.8 % under a second rule. The optimization model focuses on the extreme events, but considers the entire distribution in its analysis. A temporal analysis provides additional information to assist in the decision regarding the proper level of risk aversion.

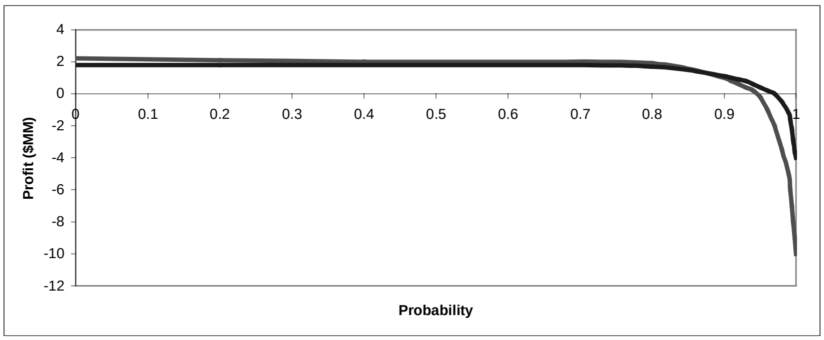

Figure 2
A Distribution of the Profit/loss before and after Optimization

1.3 Modeling Pitfalls

As with any new technology, there are pitfalls to steer clear of. We list a few common problems when implementing multi-stage planning models. Mostly, successful applications require experience with solving problems in this domain.

Difficulty #1: Overconfidence in Model's Recommendations

The recommendations of a multi-stage model should be placed in the context of the investor's background, time commitment, and ability to execute the plan. A computer model can only aid decision-making; it cannot replace sound judgment. The investor must gain an understanding of and accept the components of the system. First, the scenario generator must be understandable and consistent with historical economic relationships. Bond returns, for example, should be directly linked to changes in the underlying interest rates. Second, the decision simulator must be implementable, in the context of the investor's environment.

Optimization applies after one has developed a financial simulator and become confident that the simulator provides useful information regarding the investor's problem. First simulate, then optimize (or search for improving solutions). For long-term investors, the problem involves multiple dimensions. Thus, the optimization generates sets of dominating solutions, rather than finding a single optimal solution.

Difficulty #2: Stuck in a Local Optimal Point

Many multi-stage planning problems lead to non-convex optimization models. Here, standard optimization methods such as the one supplied in Excel will only be able to locate what are called local optimal solutions. Thus, the optimization algorithm may find a local optimal solution, and the best solution – the global solution -- may not be identified. The identified efficient frontier may not be efficient, after all.

Also, the non-convexity of multi-stage financial optimization can lead to efficient frontiers with non-smooth shapes. Figure 3 depicts an example in the context of an asset and liability management (ALM) system.

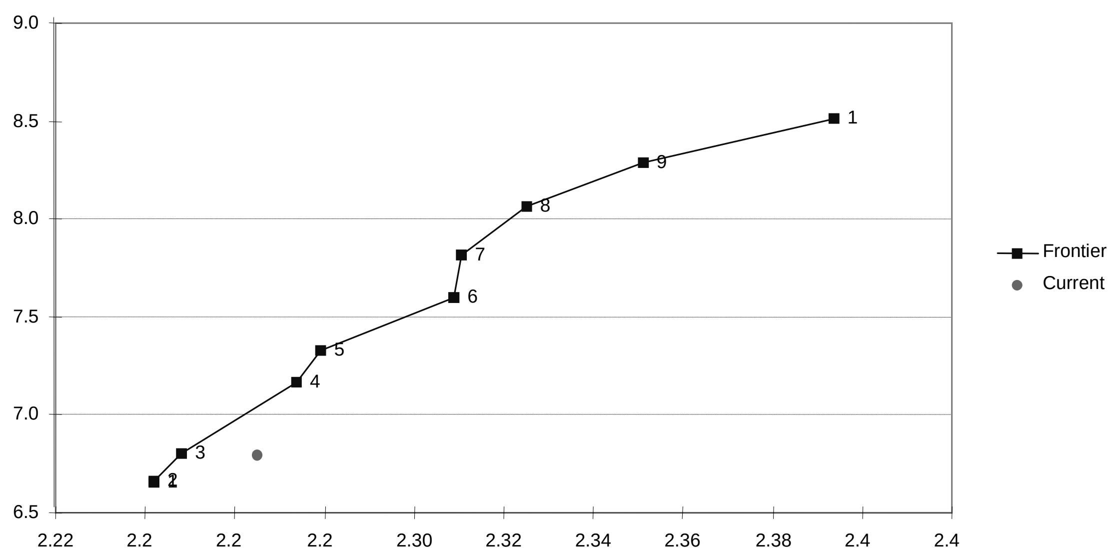

Figure 3
Example Asset-Liability Efficient Frontier for Multi-period Model
(Shape caused by non-convexity.)

Difficulty #3: Inadequate Number of Scenarios

By necessity, Monte Carlo simulations operate with a sample set of scenarios. Thus, the model's recommendations, such as the investor's expected wealth at the end of the planning horizon, will possess sampling errors. These sampling errors are akin to the errors displayed when the results of public polls are taken. Most simulation systems do not post their sampling errors. This oversight can be critical when long-tailed distributions are present and the investor is concerned about the worst-case risks. In insurance, for example, regulators are interested in the chances that the firm's capital will be adequate to pay the policyholders when future losses occur. Companies aim to protect the surplus most of the time (e.g. 99% of the years). Sampling errors can be relatively large if the system does not employ advanced variance reduction methods (e.g. Rush et al. 2000). In any event, sampling errors should be supplied when the investor asks for them.

2. Essential Elements of Multi-stage Models

Several ingredients are necessary in order to construct a multi-stage financial planning system. First, a financial simulation must combine a system (set of stochastic equations) for generating the random variables with a decision simulator (Figure 4). After these elements are successfully implemented, we can search for improving and non-dominated solutions by means of optimization algorithms. To keep the discussion short, we can only summarize the issues.

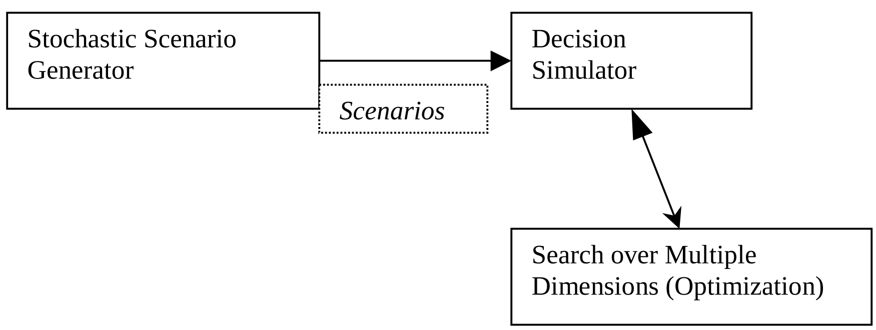

Figure 4
Elements of Multi-stage Financial Planning

A financial planning system requires a set of discrete time periods, t = {1,2,...,T}, where the planning period [1,T] encompasses the investor's important goals and obligations. Once a decision simulator is implemented and accepted[4], the investor can test out alternative investment and business (or savings) choices in order to find a suitable set of policy actions.

2.1 Scenario Generators

Financial simulation requires a systematic method for projecting the random variables over the planning horizon. We call these systems of stochastic equations by the term "scenario generators". Numerous scenario generators are in active use by actuaries, financial planners and insurance executives. A significant economic generator in the U.K. is the Wilkie model (1987). Mulvey developed a system for Pacific Mutual (Mulvey 1989) and latter for Towers Perrin –

[4] Acceptance requires a number of steps, such as backtesting, performance monitoring, and stress testing.

Tillinghast (Mulvey 1996). These generators possess a cascade structure with nonlinear relationships; hence calibration approaches such as maximum likelihood are inappropriate. Alternatively, firms have developed linear systems such as vector auto-regressive models (Cariño 1994). Linear models, while relatively easy to calibrate, can lead to possibly non-intuitive results over long-term planning periods.

The scenario generator provides a foundation for the financial simulator and latter the optimizer. Thus, the investor should pay close attention to the scenarios generated. Some questions to ask: Are the resulting scenario paths sensible? See Figure 5 for a graph of interest rates for ten scenarios. Do the asset returns link to the underlying economic factors in a consistent fashion? For example, are bond returns derived from interest rate changes, especially the spot rate and yield curves? What are the summary statistics of the generated scenarios? Does the model include an adequate number of scenarios with tail events, such as recessions?

Stochastic scenario generators are becoming established in other domains. An important example involves estimating the probability of loss under projected catastrophic events – hurricanes and earthquakes. These stochastic systems play a critical function when insurance companies and state regulators set pricing decisions. In addition, running these models helps in calculating a company's capital requirements. Many insurance companies routinely evaluate their books of business with the Cat models from AIR, RMS, and EQE international, among others. Typically, 10,000 to 100,000$^+$ scenarios are generated due to the large number of possible events and the resulting severe consequences for losses. We employ one of these systems in the next subsection.

Another area for scenario generation entails weather projections and related topics such as degree-days above or below average for selected locations and time-periods. This area will grow along with the new weather related securities and derivatives.

All of these areas require attention to three critical issues: 1) the realism of the model equations; 2) calibration of the parameters; and 3) sampling procedures. In addition, the accuracy of the projection systems should be evaluated with historical data (backtesting), as well as on an ongoing basis. Confidence in these systems will grow only if they display a sufficient degree of accuracy. Research is needed in this regard, especially with regard to points 1 and 2 above[5].

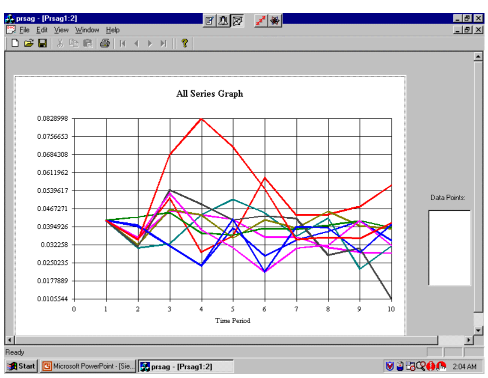

Figure 5
Illustration of 10 Interest Rate Scenarios

---

[5] Researchers are searching for improved approaches for modeling the economic environment. For example, see the paper by Tokat et al. (2000).

2.2 Decision Simulators

Financial simulators mimic the investor's decisions over the planning period. An insurance company, for example, makes decisions regarding their asset mix, business strategies, and the firm's capital structure (leverage, exposure, etc.). In a similar fashion, a pension plan administrator sets investment strategies along with pension policy for contributions, benefits and expenses. For individuals, decisions involve savings and expenditures with possible tax implications. The investment process consists of t = {1,2,3, .., T} time stages. The first decision juncture represents the current date. The end of the period, T, is called the planning horizon. Typically, it depicts a point in which the investor has some critical planning purpose, such as the repayment date of a substantial liability.

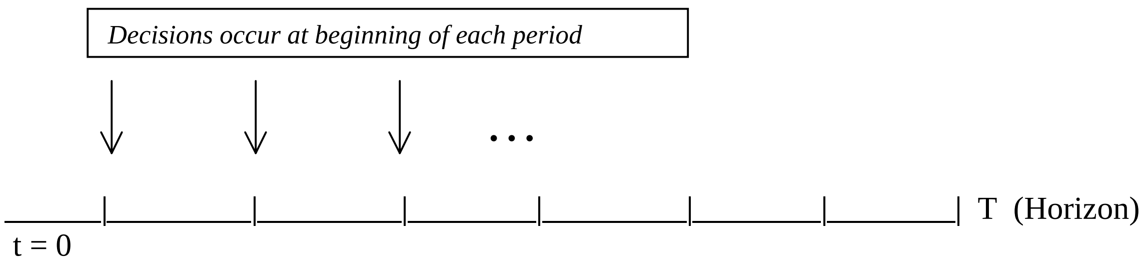

Figure 6
The Planning Period (t = 1,2, …, T)

The primary decision variables designate asset proportions, liability related decisions, and cash flow decisions:

$x_{j,t}^s$	investment in asset j, time t, scenario s
$y_{k,t}^s$	liability decision k, time t, scenario s
$u_{l,t}^s$	other cash flow l, time t, scenario s.

At each time period, t, the model moves money between asset categories, adjusting/paying liabilities, and paying off goals. There are numerous candidates for the objective function; see section 2.3. In addition, we impose constraints on the process such as limiting borrowing to certain ratios, addressing transactions costs whenever assets are bought or sold, or taking advantage of investment opportunities. As we will see, there are three fundamental modeling approaches for including constraints. Our goal is to find a set of dominating points, while maximizing several related temporal objective functions. Since we are dealing with uncertainty in a temporal setting, the dominating solutions, like all points, encompass a set of paths -- trajectories -- for the investor's wealth (or other measures such as surplus wealth). Ranking these paths is the subject of the next subsection.

There are two basic equations for the flow of funds at each time period:

Equation [1] for j^{th} asset category:

$$x_{j,t+1} = (x_{j,t}^s \times R_{j,s}^s) - p_{j,t}^s + q_{j,t}^s\,(1-t_j) \qquad \text{for asset j, time t, scenario s.}$$

where $R_{j,t}^s$ = return for asset j, time t, scenario s,

$p_{j,t}^s =$ sales of asset j, time t, scenario s,
$q_{j,t}^s =$ purchase of asset j, time t, scenario s,
$t_j =$ transaction costs for asset j.

Equation [2] for the cash flows:

$$x_{1,t+1} = (x_{1,t}^s \times R_{1,s}^s) - \sum_j q_{j,t}^s + \sum_j p_{j,t}^s (1 - t_j) + w_t^s - \sum_k y_{k,t}^s - \sum_l u_{l,t}^s$$

where w_t^s = cash inflows at time t, scenario s,
cash is asset category 1

Other linear equations are needed to model constraints, policy rules, and so on (for examples, see Ziemba and Mulvey 1998).

The multi-stage investment model must avoid looking into the future in an inappropriate fashion. The system cannot optimize over scenarios that do not represent a range of plausible outcome for the future. To prevent this occurrence, we add special constraints to the model, called non-anticipatory conditions. The general form of the constraints is:

$$x_{j,t}^{s1} = x_{j,t}^{s2}$$

for all scenarios s1 and s2 which inherent a common past up to time t (scenarios sharing a common path in Figure 8 up to t).

The financial planning system addresses these non-anticipatory conditions, either explicitly or implicitly; special purpose algorithms are available (Birge and Louveaux 1997, Dantzig and Infanger 1993, Dempster 1998, Infanger 1993, Mulvey and Ruszczynski 1995).

Next, we present an example of a real-world company-wide simulation. The problem involves a Bermuda re-insurance company. The company takes on high-severity, low frequency catastrophic risks for insurance companies on a worldwide basis. The main issue is to evaluate two investment strategies. Currently, the company keeps most of its assets in short-term U.S. government t-bills and t-notes. We compare this investment strategy to a mixture that includes equities in the portfolio. In particular, the percentage of equity depended upon the company's ability to generate insurance business and the projected losses therein. The surplus equity rule is straightforward: Generate the profit/loss distribution for the company; making sure that there is adequate capital to support the 99% loss quantile. In other words, the company can write insurance business up to the amount that would be covered at the 99% VaR at the beginning of each quarter. If losses occur, the capital base decreases by the appropriate amount. Likewise, if the assets increase or decrease due to market fluctuations, the company's ability to write business is affected. One of the Cat modeling systems estimates insurance losses; the CAP: Link system projects economic scenarios (Mulvey et al. 2000).

The "surplus-equity" strategy maintains the required capital in short-term U.S. government obligations (as before), but the remaining assets are placed in an S&P 500 index fund.

On the business side, we assume that the amount of business depends upon the severity of losses for the past several years. If there is a large Cat event (suitably defined), the insurance prices will increase, along with profits -- provides that the company has adequate capital to continue after the loss. Otherwise, if no severe events occur, the profitability of the insurance lines will slowly decrease – about 10% per year. The company's simulated response to less profitable business is to reduce the amount of business written. Given these policy rules, we simulate the company over the next five years. We assume quarterly time steps, comparing the short-term cash/notes asset strategy, with the surplus equity strategy.

The results of the simulation for 50,000 scenarios appear in Figure 7. Here, we plot profit/loss exceedance curves for both strategies – all cash/notes versus surplus equity. The surplus equity strategy stochastically dominates the cash strategy at the end of the 5-year period. However, the surplus equity strategy has slightly greater risks for intermediate years. Also, there might be concern for the correlation between the most severe earthquakes and the health of the equity market in certain countries. We assume independence between Cat scenarios and the economic scenarios. If investors are worried about this possibility, we suggest that they purchase one of the dual trigger options[6] now available on the market, or structure a capital infusion under exceptional loss events. The multi-stage simulator illustrates the risk/reward issues over a temporal setting.

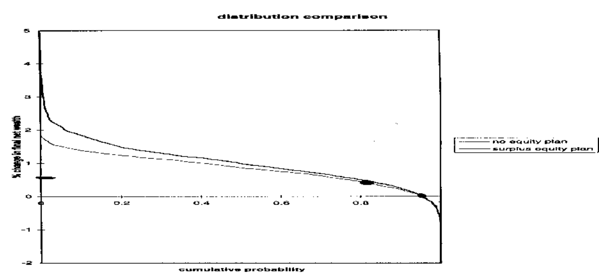

Figure 7
Profit/loss Exceedance Curve for Cash and Surplus-equity Investment Strategies
(Surplus-equity stochastically dominates cash/notes.)

2.3 Multi-stage Optimization (Multi-dimensional Search)

The optimization of a multi-stage financial planning system requires tradeoffs and compromises. Thus, the notion of a single optimal point, a.k.a. the best risk adjusted profit or maximum expected utility[7] solution, rings hollow. Instead, the optimization performs a series of searches for dominating solutions. A critical aspect entails the interface with the investor.

For a long-term investor, the search for a compromise solution is complicated by the nature of the problem. There are numerous preference metrics. An individual may wish to maximize her chance of retiring with a certain degree of comfort (say 70% of salary). Or a pension plan administrator may wish to minimize the present value of contribution over the next 10 years. The temporal aspects of these problems give rise to multiple dimensions and the inevitable tradeoffs

[6] An option that pays off when both the stock market drops AND there is a large CAT event.
[7] The original von Neumann Morgenstern utility approach did not address temporal issues. Others, including Koopmans, have extended VM utility (Keeney and Raiffa 1993) but considerable controversy remains.

that occur. The usual tension occurs between short-term risks versus long-term gains; for instance, the superior growth of equities as compared with higher volatilities over short-time periods. But there are other issues to analyze: 1) the distribution of surplus at horizon and at end of 1-year; 2) the expected contribution over time and worst case; and 3) pension expenses over time. A series of efficient frontiers can be developed for each of these pairs of dimensions.

From modeling and algorithmic perspectives, there are three fundamental approaches for solving multi-stage stochastic optimization problems[8]. The first is dynamic stochastic control, as developed by Samuelson, Merton (1992), and many others (e.g. Davis and Norman 1990 and Tasker el al. 1988). The basic idea is to construct the model so that the optimality conditions can be well posed via a system of equations. Typically, the optimality system is solved by means of discretization methods or dynamic programming rollback procedures. Several significant problems have been "solved" with these approaches, such as the investor optimal investing/saving problem (see Merton 1992, and Duffie 1992).

In many cases, unfortunately, stochastic control cannot address critical features. For instance, the liability structure of a pension plan or an insurance company depends upon a complex set of regulations, policy rules and other non-continuous functions and relationships. General constraints are often required – the sum of the equity assets must be no more than x % of the sum of the bond assets, for instance. In addition, the state space expands as a function of the number of driving economic factors and the investor's environment. As a consequence, there have been few attempts to model realistic long-term planning problems with stochastic control. An exception is Brennan and Schwartz (1998).

The second fundamental approach is multi-stage stochastic programming. This method depends upon a scenario tree as shown in Figure 8. Decisions are defined for each time period and for each scenario branch in the tree. As an example, an asset and liability system would have variables for each asset category – $x_{j,t}^{s}$ for each scenario s, for each time period t, and for each asset j. While the number of variables quickly expands, depending upon the length of the planning period, there are efficient methods for reducing problem size by bounding (Birge and Louveaux 1997) or by careful scenario sampling (Dempster 1998). It is important to emphasize that the scenario tree is not equivalent to the investor's state space. Rather, the scenarios represent samples of the random variables. As a consequence, it is straightforward to include realistic features, such as turnover constraints, market-impact costs, and complex liability structures, among others. Frank Russell conducted the first large-scale implementation of stochastic programming in financial planning for the Yasuda insurance company (Cariño, et al. 1994).

To simplify the stochastic program and greatly reduce the number of decision variables, the modeler can impose policy or decision rules. For example, the fixed-mix rule is widely employed for asset allocation (Mulvey, et al. 2000) since it is simple to understand, has desirable long-term properties, and can be optimized within the context of the multi-dimensional nature of the planning problem. We describe a pension plan example with this policy rule in the next section. Other policy rules include constant proportional portfolio insurance (Perold and Sharpe 1988), surplus-equity rules for pension plans and insurance companies, among others.

---

[8] In this paper, we do not focus on approaches that solve a sequence of single-period models since they assume that there are no market impact costs and that uncertainties display temporal independence. The interested reader should consult the introduction in Ziemba and Mulvey (1998) for further details relating to the pros/cons of alternative modeling formulations for stochastic optimization.

Once the investor chooses a policy rule, the stochastic optimization problem takes on the complexion of a Monte Carlo simulation program, with an optimization overlay. The financial simulator evaluates selected sets of policy rules and particular parameter settings. For instance, a pension planner might wish to start with standard benchmarks such as {40/60, 60/40 and 80/20} stock/bond mixes. The search for a desirable compromise solution proceeds from these benchmarks. Unfortunately, applying standard search algorithms can lead to sub-optimal (dominated) solutions due to the presence of non-convexities (Maranas, et al. 1997).

A policy rule can be interpreted as a series of constraints on the underlying stochastic program. Thus, we might expect that the recommendations of the policy rules would be worse than the stochastic program. Regrettably, the problem is complicated by several factors: 1) the multi-dimensional aspects of the choices; 2) the added complexity of the stochastic program; 3) the difficulty in computing sampling errors; and 4) computer resources needed to solve a large stochastic program.

We have found that policy optimization is the ideal approach for many long-term investors, especially those who are new to the area of multi-stage financial planning. The complexity of stochastic programming and the severe assumptions of stochastic control prevent their widespread use. On the other hand, policy optimization is readily understood, can be easily overridden by the investor, and provides sampling errors and stress test results. Of course, advanced users may wish to employ the more powerful stochastic programming or stochastic control tools. These models can serve as adjuncts to the policy rules. For example, they can assist with the search for the good re-balancing decisions. At this time, we believe that no single approach has a clear advantage over the others. It seems likely that a combination of approaches may eventually succeed, especially as computers become more powerful and easier to parallelize, and as multi-stage planning systems become more widely available.

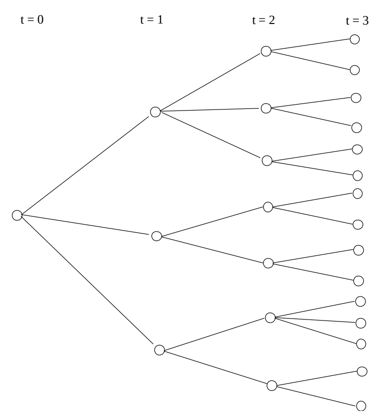

Figure 8
Sample Scenario Tree for Stochastic Programs
(Decisions are rendered at each node of this tree.)

3. Pension Plan Example

In this section, we describe the highlights of a multi-stage analysis for a large pension plan[9]. The multi-stage model provides information that is difficult to gather in a single period approach. Significant issues such as rebalancing rules and temporal tradeoffs can be evaluated.

Carrying out an ALM study requires a set of capital market assumptions. After discussions with the company executives, we selected a 10-year planning horizon, and eleven asset categories, along with the corresponding projected expected returns and annual volatilities. Table 1 presents the relevant information for the 10-year planning period. These projections combine historical values (e.g. volatility), with judgmental factors (e.g. equity risk premiums) and current market conditions (e.g. long bond rates). In addition, a covariance matrix was generated in a similar manner, mostly with historical values for the correlations. Due to policy rules, several asset categories were given upper limits, such as private equity at 6%.

Given this data, it is straightforward to generate a single-period Markowitz efficient frontier for the assets (Figure 9). The lower-risk side of the frontier contains low volatility assets (e.g. cash and bonds), while the upper-risk side contains mostly equity. Remember that the top of the frontier is a linear program; thus, the highest returning asset is emphasized to the degree possible. There is little or no diversification, unless the investor puts upper bounds on higher returning assets. The client’s current position is in the moderate risk area. Note that the highest returns lie between 8.5 % and 9 % due to policy constraints on equity assets.

Table 1
Capital Market Assumptions for Pension Plan Example

Asset category	Expected Returns	Annual Volatility
Cash	5 %	1 %
S&P 500 equity	9 %	16 %
Private equity	14 %	24 %
European equity	9 %	18 %
Far east equity	9 %	18 %
Emerging markets	9 %	27 %
Corp/Govt bonds[10]	7 %	6 %
International bonds	6.8 %	7 %
High-yield bonds	8.3 %	8 %
Real estate	7.8 %	14 %
Hedge funds	8.3 %	7 %

[9] This example is patterned after a real world pension plan, but the numbers have been modified and identities disguised.
[10] Salomon BIG index

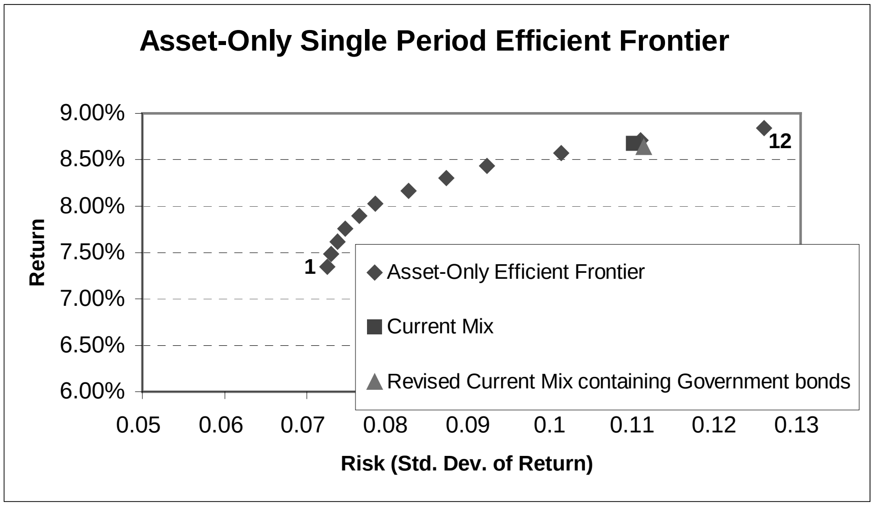

Figure 9
Single-Period Asset Frontier

As the next step, we developed a multi-stage scenario generator for modeling the underlying economic factors, accompanied by the corresponding returns for the eleven assets. The factors consisted of interest rates, inflation, economic activity, and corporate earnings in the US. A global model was deemed inappropriate due to the nature of the company. The calibration targets for asset returns are identical to the single-period model. In addition, however, we required temporal series for the underlying factors so that the liabilities could be generated in a consistent fashion with the assets. The calibration was successful; the algorithm generalizes Duffie and Singleton's simulated moment estimation (Berger and Madsen 1999).

Based on these scenarios, we solved a non-convex optimization model for determining the best asset mixes over the 10-year planning horizon. The asset efficient frontier at the end of the planning horizon provides recommendations that are different from the single-period model. First, the top of the efficient frontier has a mixture of equity assets, rather than the typical single asset coming out of a single-period optimizer. The optimal solution in terms of returns encompasses a range of equity assets, including S&P500, European and Far East, private equity, and emerging market equity. The investor improves performance by diversifying and rebalancing at each period to a target proportion.

Also, the returns are higher – almost 9.6 % at the maximum. Higher returns are due to the re-balancing gains attributed to the fixed-mix policy rule[11]. The multi-period model diversifies the portfolio to a much greater degree than the single-period model, especially for the moderate-to-high risk tolerant investors[12].

[11] We re-balanced the portfolio each quarter.
[12] Long-term investors should be able to take on higher short term risks, especially if they can withstand reasonable levels of volatility induced by recessions, short-term crises and related events.

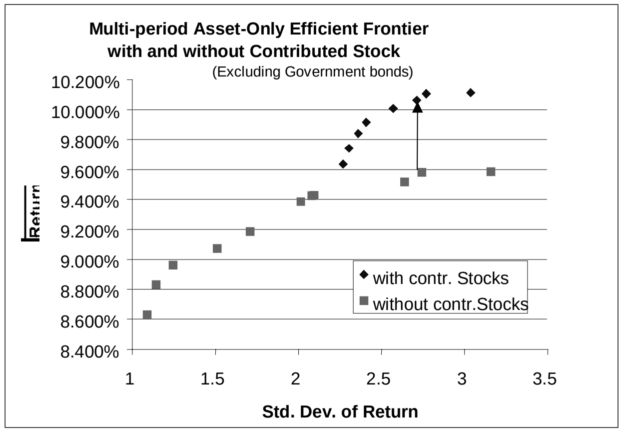

Figure 10
Multi-period Asset Efficient Frontier
(Returns are higher than single-period model due to re-balancing gains.)

Once a multi-stage scenario generator/decision simulator is constructed, the investor can readily evaluate a number of issues. To give an example, we test the impact of adding contributed stock from two subsidiaries (Table 2). The contributed stocks will be sold over the next 3 to 5 years in a linear fashion and replaced by S&P 500 stock. The goal is to increase the company's surplus position so that it can better withstand difficult economic conditions in the future. Not only does the contributed stock improve the pension surplus, but also it enhances returns at the higher ends of the frontier (Figure 10), despite the fact that the contributed stocks possess the usual 9 % equity expected returns and will be sold within 3-5 years. Greater returns are due to re-balancing gains; the contributed stock's volatility (39 and 44 %) gives an advantage to long-term investors. See Luenberger (1998) for an explanation of re-balancing gains[13] and the role of volatile assets in multi-stage models.

Table 2
Additional Asset Categories

Asset category	Expected Returns	Annual Volatility
Contributed stock 1	9 %	44 %
Contributed stock 2	9 %	39 %
US government bonds	6.75 %	10 %

[13] Luenberger calls the strategy "volatility pumping." Market impact costs will dampen these results, and a more comprehensive analysis is needed to find suitable re-balancing strategies.

The recommendations derived from a 10-year efficient frontier must be tempered by the analysis of shorter-term impacts. Can the company tolerate the year-to-year uncertainties? There are many ways to evaluate this issue. For instance, we are interested in the plan's contributions and expenses over time. There distributions are easily computed from the ALM system.

Last, we show the optimal asset solution from the standpoint of a surplus analysis. Here, we simulate the actuarial and accounting requirements for paying beneficiaries, for making contributions and paying PBCG expenses. The company's actuaries supplied the data so that we could estimate these values for each scenario over the 10-year planning horizon. The full pension plan simulator can then be optimized in order to locate sets of dominating solutions. As mentioned, the actual problem consists of multi-dimensions; it cannot be simplified to a single or even a small number of graphs. Never the less, we will depict two surplus efficient frontiers in order to emphasize an important point – namely, the suitability of asset categories should be fixed in concert with their relationship to the liabilities. To illustrate, we add a long-government bond asset category to the analysis. This asset more closely links to the economic changes in the liabilities than any of the other assets. Thus, it gives greater protection during economic downturns in which interest rates drop (Figure 11) than other bond categories. In addition, the long-bond asset can improve returns over longer periods of time due to its volatility[14]. The 10-year surplus efficient frontier (Figure 12) also improves by the addition of government long-bonds. Again, the multi-stage model gives insights into the nature of long-term investment strategies.

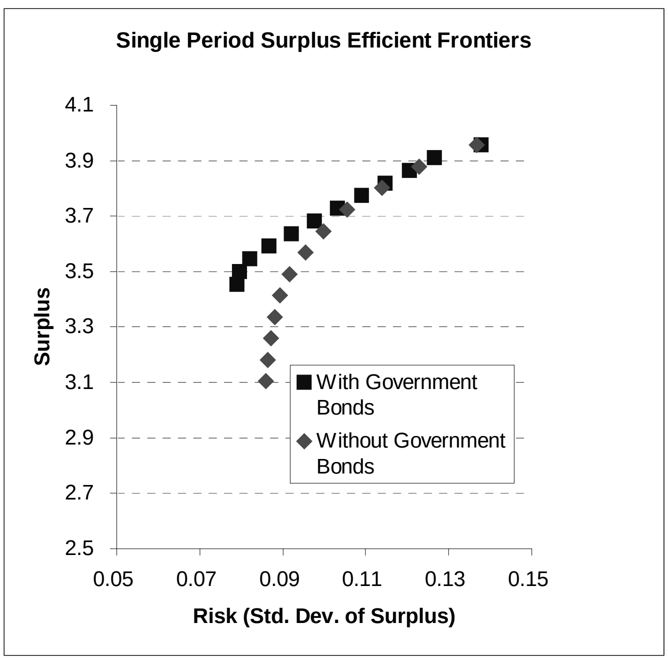

Figure 11
Surplus Efficient Frontier at End of Year One
(Government bonds provide cushion against economic downturn.)

[14] The expected returns from long bonds are slightly less than the Salomon BIG, but have greater volatility. Thus, long-bonds take advantage of volatility pumping over long time horizons.

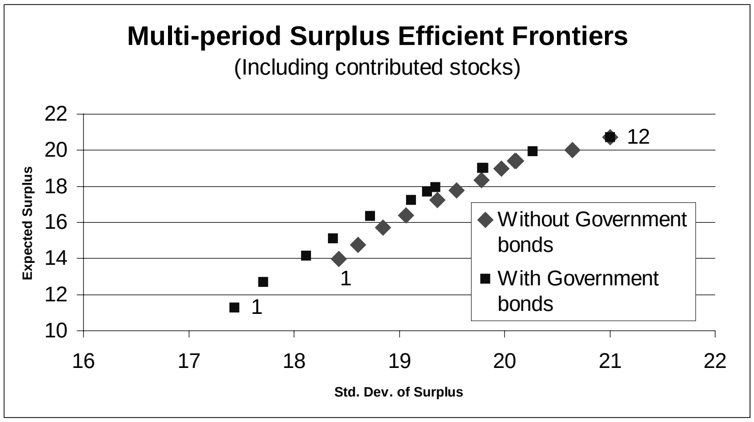

Figure 12
Surplus Efficient Frontier at end of 10-Years
(Government bonds take advantage of volatility pumping.)

4. Conclusions

Multi-stage simulation/optimization models provide advantages over single-period approaches. Not only can the models include important details such as transaction and market impact costs, but also careful dynamic investment strategies can improve performance, especially for long-term investors. The situation of re-balancing gains, section 3, is an example. Here, the investor takes advantage of volatility by maintaining a target fraction of wealth in each asset category (give or take specified ranges). A more sophisticated version of this strategy might take advantage of other phenomena, e.g. the natural mean reversion of volatility itself.

What are barriers to wider use of the multi-stage technology? First, investors should become aware of the feasibility of optimizing multi-stage financial planning models. Modern computers are able to handle the complex task of solving these large, possibly non-convex, optimization problems. Second, the scenario generators will need to become more fully tested, especially for the rare events. In some cases, there is insufficient data to make an informed projection, such as the expected returns of the volatile high technology sector over an extended time horizon. There is little data on which to base these estimates. Research is needed for modeling the extreme events, especially in the area of conditionality across securities.

What are other topics for future research? There is need for mechanisms for helping long-term investors resolve the multi-dimensional and temporal aspects of investment decision-making. The interface environment is critical in this regard. Second, research is needed in combining the best aspects of the various approaches to optimization – feedback between the concepts, for example, and establishing bounds on the results of policy optimization. Better understanding of the pros/cons of multi-stage financial planning systems – both simulation and optimization-based models – for selected financial planning problems will be required.

In the end, we can turn to engineering for ideas concerning the proper level of risk protection. For instance, aeronautical engineers optimize the shape of modern airplanes by minimizing weight and maximizing strength. Computer simulation/optimization systems are employed. The Boeing 777, for instance, was largely designed with the aid of computer analysis. An important part of this analysis is stress testing the airplane in conjunction with severe conditions (Figure 13). The computer flies the airplane through these scenarios in order to evaluate the system's ability to withstand the stresses. In a financial simulation, we seek similar goals – to evaluate the investor's ability to withstand economic downturns. For example, a U.S. pension plan surplus will likely drop during a recession with lower equity returns and lower interest rates. The simulator will give information regarding the accompanying contributions. A long-term investor should be able to withstand these periods. Multi-stage financial planning systems are helpful in gaining an understanding of future risks and the investor's planned responses.

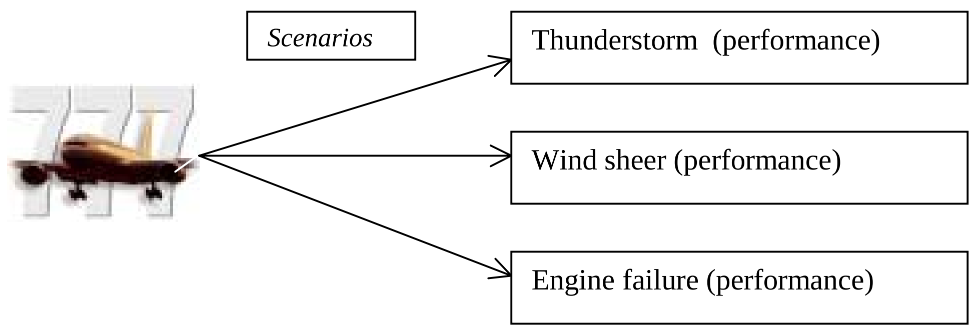

Figure 13
Computer Simulation of an Airplane
(Engineering designs are routinely stress tested by scenario analysis.)

References

A. Berger and C. Madsen, "A Comprehensive System for Selecting and Evaluating DFA Model Parameters," Casualty Actuarial Forum, summer 1999.

A. Berger and J. Mulvey, "Integrative Risk Management for Individual Investors," in **Worldwide Asset and Liability Modeling**, (eds. W. Ziemba and J. Mulvey), Cambridge University Press, 1998.

J. Birge and F. Louveaux, **Introduction to Stochastic Programming**, Springer, 1997.

M. Brennan and E. Schwartz, “The Use of Treasury Bill Futures in Strategic Asset Allocation Programs,” in **Worldwide Asset and Liability Modeling**, (eds. W. Ziemba and J. Mulvey), Cambridge University Press, 1998

D. Cariño, T. Kent, D. Myers, C. Stacy, M. Sylvanus, A. Turner, K. Watanabe and W. Ziemba, “The Russell-Yasuda Kasai Model: An Asset Liability Model for a Japanese Insurance Company using Multi-stage Stochastic Programming,” *Interfaces*, **24**, 29-49, 1994.

G. Dantzig and G. Infanger, “Multi-stage Stochastic Linear Programs for Portfolio Optimization,” *Annals of Operations Research*, **45**, 59-76, 1993.

M. Davis and A. Norman, “Portfolio Selection with Transaction Costs,” *Mathematics of Operations Research*, **15**, 676-713, 1990.

M. Dempster, “The Development of the Midas Debt Management System,” in **Worldwide Asset and Liability Modeling**, (eds., W.T. Ziemba and J. M. Mulvey), Cambridge University Press, 1998.

D. Duffie, **Dynamic Asset Pricing Theory**, Princeton University Press, 1992.

D. Duffie and K. Singleton, “Simulated Moment Estimation of Markov Models of Asset Prices,” *Econometrica*, **61**, 929-952, 1993.

G. Infanger, **Planning under Uncertainty: Solving Large-Scale Stochastic Linear Programs**, Boyd and Fraser, 1993.

R. Keeney and H. Raiffa, **Decisions with Multiple Objectives: Preferences and Value Tradeoffs**, John Wiley, 1976. Reprinted by Cambridge University Press, Cambridge, MA., 1993.

D. Luenberger, **Investment Science**, Oxford University Press, 1998.

C. Maranas, I. Androulakis, C. Floudas, A. Berger and J. Mulvey, “Solving Stochastic Control Problems in Finance via Global Optimization,” *Journal of Economics Dynamics and Control*, **21**, 1405-1425, 1997.

H. Markowitz, **Portfolio Selection: Efficient Diversification of Investments**, Second Edition, Blackwell, 1991.

R. Merton, **Continuous Time Finance**, Blackwell, 1992.

J. Mulvey, “A Surplus Optimization Perspective,” *Investment Management Review*, **3**, 31-39, 1989.

J. Mulvey and W. Ziemba, “Asset and Liability Allocation in a Global Environment,” in **Finance**, R. Jarrow *et al.*, Eds., *Handbooks in Operations Research and Management Science*, Vol. 9, North Holland, 435-463, 1995.

J. Mulvey and A. Ruszczynski, “A New Scenario Decomposition Method for Large-Scale Stochastic Optimization,” *Operations Research*, **43**, 477-490, 1995.

J. Mulvey, “Generating Scenarios for the Towers Perrin Investment System,” *Interfaces*, March 1996.

J. Mulvey, G. Gould and C. Morgan, “An Asset and Liability Management System for Towers Perrin – Tillinghast,” *Interfaces*, **30**, 96-114, 2000.

A. Perold and W. Sharpe, “Dynamic Strategies for Asset Allocation,” *Financial Analysts Journal*, 16-27, Jan. 1988.

R. Rush, J. Mulvey, J. Mitchell and T. Willemain, “Stratified Filtered Sampling in Stochastic Optimization,” *J. of Applied Mathematics and Decision Sciences*, **4**, 17-38, 2000.

M. Taksar, M. Klass and D. Assaf, “A Diffusion Model for Optimal Portfolio Selection in the Presence of Brokerage Fees,” *Mathematics of Operations Research*, **13**, 277-294, 1988.

Y. Tokat, S. Rachev and E. Schwartz, “The Stable non-Gaussian Asset Allocation: A Comparison with the Classical Gaussian Approach,” presented at the International Symposium on Mathematical Programming, Atlanta, summer, 2000.

A. Wilkie, “Stochastic Investment Models: Theory and Applications,” *Insurance: Mathematics and Economics*, **6**, 65-83, 1987.

W. Ziemba and J. Mulvey, (editors), **Worldwide Asset and Liability Modeling**, Cambridge University Press, 1998.