Calibration of Stochastic Scenario Generators for DFA

John M. Mulvey, Ph.D.,
François Morin, FCAS, MAAA, and
Bill Pauling

Calibration of Stochastic Scenario Generators for DFA

John M. Mulvey, Ph.D.
Princeton University
Princeton, NJ 08544
Email: mulvey@princeton.edu

François Morin, FCAS, MAAA
Tillinghast – Towers Perrin
Hartford, CT
Email: Morinf@towers.com

Bill Pauling
Towers Perrin
Philadelphia, PA
Email: paulinb@towers.com

April 1999

Abstract

Leading actuarial companies employ stochastic simulation models to evaluate the viability of pension plans and insurance companies over a set of projected scenarios. A critical element involves generating future scenarios. We show that the problem of calibrating a stochastic scenario system can be posed as a special optimization model, and illustrate the process by means of the Towers Perrin – Tillinghast CAP:Link system. We briefly discuss solution algorithms for the resulting non-convex problem. Areas for future research are indicated.

I. Introduction to Integrated Financial Risk Management

Over the past several years, innovative insurance companies have begun building integrated financial risk management systems. These efforts aim to evaluate the company's activities within a common framework. The goal is to maximize shareholder wealth by focusing on the overall risks and rewards to the organization as measured in several ways. Ideally, the major areas affecting the company's results should be integrated: asset allocation, business management, corporate structure, and re-insurance. Doing so provides the best opportunity to achieve the company's goals over time.

For some insurance businesses, the importance of linking asset and liability risks is well understood. An example is an annuity whose payoff is set at a proportion of the US S&P500 stock return above a set index. It would be foolhardy to invest the assets for this business without understanding the risks of mismatching assets and liabilities. Given recent performance, it is a simple matter to assign assets that exactly match the product's payoff patterns, through options or futures contracts, or funding reserves by dynamically purchasing or selling the stock. Asset and liability management for this type of insurance is a clear and obvious concern. For all insurance companies, there is need to manage the assets and liabilities so that surplus will grow at a rapid pace, as compared with maintaining the surplus at a constant or slowly growing trajectory. In addition, shareholders will seek out insurance companies that grow rapidly and possess diversification benefits.

A dynamic financial analysis (DFA) system consists of three major elements: a stochastic scenario generator, a multi-period simulator, and an optimization module (Figure 1). The first two elements form the corporate simulation system; these are deployed before the optimization module searches for the best compromise decisions given the relevant business, policy, and regulatory constraints. In effect, the optimization runs the simulation by identifying the combination of decisions that best fits the proposed objective function over the multi-period planning horizon.

A critical issue involves constructing the economic scenarios. Each scenario depicts a single coherent path for the primary uncertainties, such as interest rates, inflation, and business activity. Typically, the scenarios are generated by sampling from a system of stochastic differential or difference equations. As a simple example, we could generate short government interest rates by means of a mean reverting equation:

$$dr_t = a\,(r_\mu - r_t)\,dt + \sigma\,dZ \qquad (1.1)$$

where dZ = Wiener white noise term
r_t = interest rates at time t.

Here, the equation shows that the change in interest rates at time t depend upon three factors – the distance to the target reversion parameter (r_μ), the drift parameter (a), and the instantaneous volatility (σ). Thus, there are three parameters associated with equation (1.1). These parameters dictate the characteristics of the sample paths[1]. The calibration process determines the appropriate values for these parameters. We call the approach – integrated parameter estimation (IPE). The basics are taken up in the next section.

[1] Three sources of errors must be considered in a DFA system: model error, calibration error, and sampling error. We are solely concerned with the second source in this paper. See Mulvey and Madsen (1999) for a further discussion of addressing errors in DFA systems.

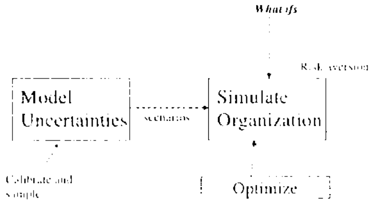

Figure 1
Components of DFA Technology

2. An Optimal Fitting Model for Calibration

This section describes the calibration of the scenario generator as a special optimization model. The primary concept is to match summary statistics and other indicators such as inter-quartile ranges and quantiles as closely as possible, while defining the model parameters as decision variables in the optimization model. The approach directly traces to traditional fitting models, including maximum likelihood, method of moments, and simulated moment estimation. As with these approaches, the model parameters are determined by reference to specialized optimal fitting problems.

Judgement is necessary when determining the parameters of a stochastic model. Fixing parameters is equivalent to setting assumptions. Ideally, we would test the impact of various settings of the parameters on the model's recommendations as shown in Figure 2.

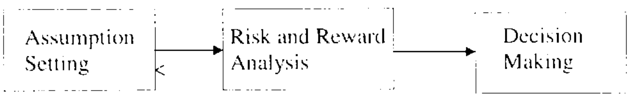

Figure 2: The Three Stages for Conducting a DFA

The process entails combining feedback and revision in order to become comfortable with the recommendations.

2.1 The Estimation Problem

This section reviews the generalized method of moments (GMM) of Hansen (1982) and the simulated moments estimation (SME) of Duffie and Singleton (1993). The notation follows Duffie and Singleton (1993).

Consider a function $H: \Re^N \times \Re^P \times \Theta \to \Re^N$, with parameter set $\Theta \subset \Re^Q$. Also consider an observation function $f: \Re^{NL} \times \Theta \to \Re^M$. A scenario process $\{Y_t\}_{t=1}^{\infty}$ is generated by the difference equation

$$Y_{t+1} = H(Y_t, \varepsilon_{t+1}, \beta_0) \tag{2.1}$$

where the parameter vector β_0 is to be estimated, and $\{\varepsilon_t\}$ is an i.i.d. sequence of random variables defined on a given probability space. Let $Z_t = (Y_t, Y_{t-1}, \ldots, Y_{t-L+1})$ defining the state of the process over time. Estimation of β_0 is based on the statistics of the observed process $f_t^* \equiv f(Z_t, \beta)$. For example, we might be interested in the means and standard deviations of the asset returns.

2.1.1 Generalized Method of Moments

Define $\beta \in \Theta$ to be an arbitrary parameter setting. When $f: \beta \mapsto E\left[f(Z_t, \beta)\right]$ is analytically known and time independent, the estimation of β_0 can be done with the generalized method of moments (GMM). For these cases, the estimator is:

$$b_T = \arg\min_{\beta \in \Theta} GMM_T(\beta)' W_T(\beta) GMM_T(\beta) \tag{2.2}$$

where $f_t^* \equiv f(Z_t, \beta_0)$, $GMM_T(\beta) = \frac{1}{T}\sum_{t=1}^{T} f_t^* - E\left[f(Z_t, \beta)\right]$, $W_T(\beta)$ is a $M \times M$ positive-definite symmetric weighting matrix, and T is the actual number of observations on f_t^*. In words, the vector b_T is the solution[2] to the minimization model defined by the least square GMM function. Hansen (1982) shows that the above minimization produces the estimator with the smallest asymptotic covariance matrix if $W_T(\beta) = \left[E\left[GMM_T(\beta) GMM_T'(\beta)\right]\right]^{-1}$.

2.1.2 Simulated Moments Estimator

For wider classes of problems where the GMM assumptions fail, the mapping $f: \beta \mapsto E\left[f(Z_t, \beta)\right]$ may be replaced by its simulated version. We assume a sequence $\{\hat{\varepsilon}_t\}$ of

[2] The arg min notation refers to the solution of the posed optimization model.

random variables with a distribution identical to and independent of $\{\varepsilon_t\}$. The simulated state process $\{Y_t^\beta\}$ occurs by choosing a starting point Y_1^β, and letting

$$Y_{t+1}^\beta = H(Y_t^\beta, \hat{\varepsilon}_{t+1}, \beta) \quad t = 1, 2, \ldots \tag{2.3}$$

The simulated observations (summary statistics) are defined as $f_t^\beta = f(Z_t^\beta, \beta)$. If $\{f_t\}$ and $\{f_t^\beta\}$ are governed by the law of large numbers, and identification conditions (Duffie and Singleton (1993)) are met then, $\lim_{T \to \infty} SME_T(\beta) = 0 \Leftrightarrow \beta = \beta_0$.

The SME estimator is,

$$b_T = \arg \min_{\beta \in \Theta} SME_T(\beta)' W_T(\beta) SME_T(\beta) \tag{2.4}$$

where $SME_T(\beta) = \frac{1}{T}\sum_{t=1}^{T} f_t - \frac{1}{\tau(T)}\sum_{t=1}^{\tau(T)} f_t^\beta$, and $\tau(T)$ is the simulation sample size for a given T. The b_T vector is the solution to the equation (2.4). Proofs of consistency (strong and weak) and asymptotic normality are given in Duffie and Singleton (1993). Additional discussions of parameter estimation through simulation can be found in Hansen and Singleton (1982), Pakes and Pollard (1989), McFadden (1989), Gregory and Smith (1990), and Lee and Ingram (1991).

2.2 Integrated Parameter Estimation

The integrated parameter estimation (IPE) approach extends simulated moments estimation in two ways. First, the target vector includes a variety of descriptive statistics besides moments, for example, serial correlation, distribution percentiles, and range estimates. Lee and Ingram (1991) allow for serial correlation in the data set, but they require that the criterion function be continuous in the mean. In contrast IPE does not require a continuous objective function; any general function or even process can be employed. Second, we place bounds on the values of selected parameters. These constraints assist in the search for the best solution.

Given a vector of parameters (β), as decision variables the IPE estimator solves the following optimization model :

$$b_T = \arg \min_{\beta \in \Theta} GSME_T(\beta)' W_T(\beta) GSME_T(\beta) \tag{2.5}$$

$$s.t. \quad L_i \le T_i - S_i \le U_i, \quad i \in \Psi$$

$$\beta_l \le \beta \le \beta_u$$

where $GSME_T(\beta) = \mathbf{m}(T_i - S_i)$, $\mathbf{m}(\cdot)$ is the IPE objective function, S_i denotes the model statistics, T_i denotes the target statistics. β_l and β_u are bounds on the parameter vector β, and Ψ indexes the pertinent statistics. When the objective function equals a distance metric, the set Ψ includes only moments, and the feasible region is unconstrained, IPE is equivalent to SME. If the weighting matrix W_T is diagonal, then (2.5) reduces to $\arg \min_{\beta \in \Theta} \sum_{i \in \Psi} \hat{W}_i [\mathbf{m}(T_i - S_i)]$, where $\hat{W}_i$ are the diagonal elements.

The IPE approach fits simulated samples from the stochastic model to a given set of descriptive statistics. Each of these descriptive statistics serves as a target, and deviations from the targets are expressed as constraints in (2.5) with tolerances L_i and U_i. The feasible region is determined by a user specified tolerance level, the maximum allowable difference between a given summary statistic and its target. One can also penalize constraints, rather than keeping them explicit. Suitable penalty functions include absolute value and quadratic functions which penalizes underage differently than overage, and asymmetric risk measure. The penalty function relatives must account for differences in scale as well as serving as importance factors. See Mulvey, Rosenbaum, and Shetty (1996) for further discussion, and Berger and Madsen (1999) for a similar approach.

The actual parameter setting process combines the actuary's judgement with the computational ability of the IPE calibration tool. Typically, several iterations are required in order to find the most desired combination of penalties and constraints to meet the target goals (equation 2.5).

3. Dealing with Non-Convexity

The solution of the calibration model is complicated by the presence of non-convexities. At its simplest, non-convexity causes standard hill climbing algorithms to stall at local optimal points. Thereby, software systems such as Microsoft's Excel solver may not find the global optimal solution. Figure 3 shows an example of numerous local optimal solutions.

The search procedure must extend itself in order to cope with non-convexities. To do this, we employ the Tabu search method, one of the most successful methods for overcoming these difficulties. The approach depends upon several memory functions that guide the search and pass through local optimal points as needed. Both long-term and short-term memory are employed.

Tabu search has proven effective for solving combinatorial optimization problems; see Glover 1990 and 1995. The procedure provides for an efficient search of a feasibility region by monitoring key attributes of the points that comprise the search history. Potential search iterates possessing attributes that are undesirable with respect to those already visited become tabu; appropriate penalties discourage the search from visiting them.

Consider a general non-convex optimization problem of the form: minimize$_x$ $f(x)$, $x \in X$. (The function $f(x)$ indicates the responses of a system to a given strategy or decision vector x.) For deterministic problems, there is a single response associated with any x. Our adaptation of tabu search has three basic elements:

1) a function $g(x) = f(x) + d(x) + t(x)$. The function $d(x)$ penalizes x for infeasibility. The function $t(x)$ penalizes x for being labeled tabu.
2) the current iterate x_c, and
3) a neighborhood of the current point N_c. The procedure generates a new iterate x_{new} by selecting the element of N_c for which $g(x)$ is smallest.

The tabu restrictions represented in $t(x)$, can address short-, intermediate-, and long-term components of the search history. Short-term monitoring is designed to prevent the search from returning to recently visited points, allowing the procedure to "climb out of valleys" associated with local minima. Short-term monitoring can also serve as a rudimentary diversification vehicle. Intermediate- and long-term monitoring techniques provide for a much more effective diversification of search over the feasible region. The $t(x)$ function in our version of tabu search relies on exploitation of short-term search history. Details of three processes are required to define our adaptation: formation of the neighborhood of the current point, assignment of tabu penalties, and termination of search procedure. See Glover, Mulvey, and Hoyland (1995).

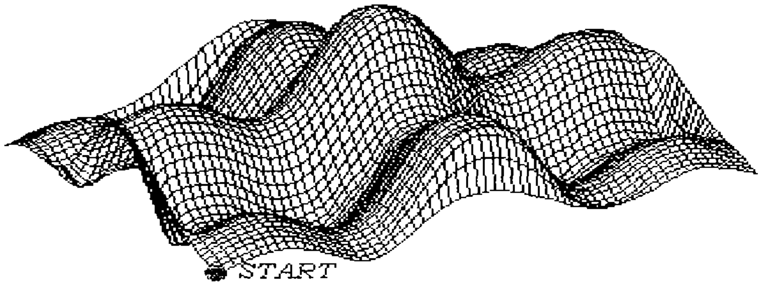

Figure 3
A Non-convex Region with Numerous Local Solutions

4. Calibration Example

The Towers Perrin – Tillinghast company employs the CAP:Link/OPT:Link system for helping pension plan and insurance clients understand the risks and opportunities related to capital market investments. The scenarios generated by CAP:Link contain key economic variables such as price and wage inflation, interest rates for twelve maturities (real and nominal), stock dividend yields and growth rates, and currency exchange rates through each year for a period of up to 20 years. We model returns on asset classes and liability projections consistent with the underlying economic factors, especially interest rates and inflation. The economic variables are simultaneously determined for multiple economies within a common global framework. Long-term asset and liability management is the primary application.

The global CAP:Link system forms a linked network of single country modules. The three major economic powers – the United States, Germany, and Japan -- occupy a central role, with the remaining countries designated as home or other countries. We assume that the other countries are affected by, but do not impact the economies of the three major countries. The basic stochastic differential equations are identical in each country, although the parameters reflect unique characteristics of each particular economy. Additional countries can be readily included in the framework by increasing the number of other countries.

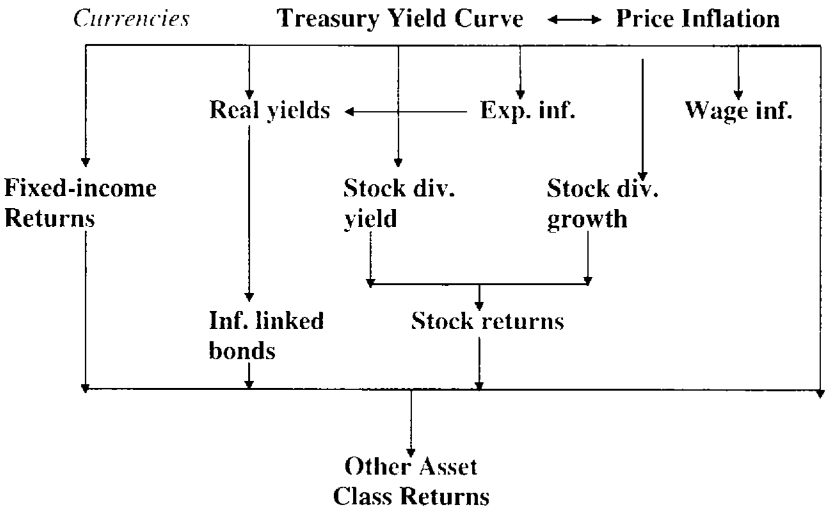

Figure 4: The cascade CAP:Link structure within a single country. Each country in Global CAP:Link depicts a common heritage.

Within each country, the basic economic structure is illustrated in Figure 4. Variables at the top of the structure influence those below, but not vice-versa. This approach eases the task of calibrating the model's parameters. The ordering does not reflect causality between economic variables, but rather captures significant co-movements. Linkages across countries occur at various levels of the model -- for example, interest rates and stock returns. These connections are discussed in Mulvey and Thorlacius (1998). Roughly, the economic conditions in a single country are more or less affected by those of its neighboring countries and by its trading partners. The degree of interaction depends upon the country under review.

The structure is based on a cascade format. Modules above and equal to that module can affect each sub-module within the system. Briefly, the first level consists of short and long interest rates, and price inflation. Interest rates are a key attribute in modeling asset returns and especially in coordinating the linkages between asset returns and liability investments. To calculate a pension plan or an insurance company's surplus, we must be able to discount the projected liability cash flows at a discount rate that is consistent with bond returns, under each scenario. Also, since dynamic relationships are essential in risk analysis, the interest rate model forms a critical element.

The second level entails real yields, currency exchange rates and wage inflation. At the third level, we focus on the components of equity returns: dividend yields and dividend growth.

Returns for the remaining asset classes form the next level, with fixed income assets reflecting the term structure of interest rates and other mechanisms. Each economic variable is projected by means of a stochastic differential equation -- relating the variable through time and with the stochastic elements of the equation and, of course, to other variables and factors at the same or higher levels in the cascade.

A critical feature for a global scenario generator is the currency model. Several complicating issues arise when modeling currency exchange rates. First, currencies must enforce the arbitrage free condition among spot exchange rates and among forward rates with differential interest rates. The second issue involves symmetry and numeraire independence: we must create a structure in which the distribution of currency returns from country A to B has the same distribution as returns from B to A. Both issues limit the form of the currency exchange models, especially when integrating three or more currencies. To avoid these problems, we focus on the strength of each country's currency. Exchange rates follow as the ratio between the strengths of any two countries. The absolute strength of any currency is a notional concept: the relative levels reflect the difference in the exchange rates. See Mulvey and Thorlacius (1998) for further details.

4.1 Example of Calibrating a Scenario Generator with both Assets and Liabilities

We now present an example of calibrating a DFA model that includes both asset and liabilities. In this example, we calibrate the CAP:Link model to produce liability growth, as well as asset returns. We then use the OPT:Link system to find a set of efficient portfolios for a hypothetical property/casualty insurer. These efficient portfolios comprise the asset-liability efficient frontier (ALEF[tm]) for the DFA. The IPE approach forms the basis for the automatic calibration tool.

4.1.1 Form of the Liability Model

For this example, we are interested in modeling a line of insurance that relates to medical and legal inflation. Liability inflation is modeled as a function of its value in the prior period, price inflation, and random volatility. The user inputs consist of an initial rate of inflation, and an assumption of future inflation. The model has two additional calibration parameters: a parameter that determines the sensitivity to modeled price inflation, and a parameter that determines the amount of random volatility. These two parameters will be calibrated in conjunction with the standard CAP:Link parameters. The Lattice optimization solver carries out the non-convex search.

4.1.2 The Calibration Process

We propose four steps for conducting a calibration exercise as shown below. It is advisable to get actuaries and users involved in the process at an early stage so that everyone understands the issues and is comfortable with the resulting model parameters.

Step 1: Analyze Historical Data

The first step to any calibration process should start with historical data. We analyzed historical data to determine the characteristics of the index. For Medical CPI we took data covering the 1947-1998 period. The historical data on Legal Services CPI is much shorter, covering the period from 1986-1998.

Step 2: Set Targets

From our analysis of historical data we determined the following targets:

	Medical CPI	Legal Services CPI
Standard deviation	1.9 – 2.2%	0.8 – 1.0%
Correlation to CPI	0.6 - .0.7	0.45 – 0.55
Average spread over CPI	0.9	0.7

We express the targets as ranges. These targets depict a blend of historical experience and forward-looking analysis. First, we start with the ranges that are consistent with historical experience. Then we adjust for historical trends. For example, for the last 11 years medical CPI has outpaced CPI by 2.5%. Can we reasonably expect this trend to continue? Over a long-term horizon, we might expect the growth in medical costs to be closer to CPI. This issue must be solved by the model developers so that proper targets can be set.

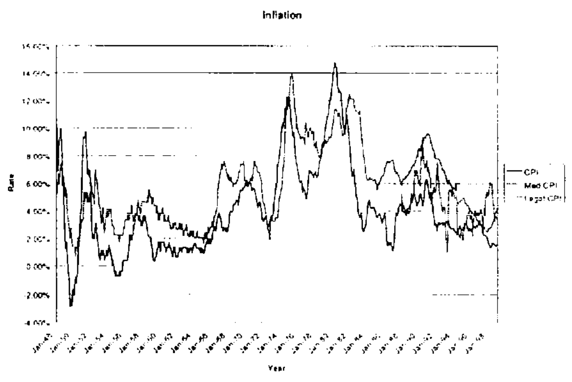

Figure 5
Historical Data for Target Inflation Series

Step 3: Use the Calibration Tool

The calibration system solves the IPE optimization model presented in section 2.2 (equation 2.5). We set up the calibration tool to run 100 scenarios per iteration. We have found through experience that 100 scenarios is a small enough number to run quickly, yet 100 scenarios produce a large enough sample to be representative of the 500 or 1,000 simulations we typically run. We have also discovered through experience that it is best to base calibrations on pure normative conditions. Calibrating to normative conditions removes the effect of trends, as initial conditions move toward their normative states. Depending on the differences between initial and normative conditions, these trends may be significant. If the trends are significant, then they may prevent the calibration tool from being able to meet the targets. Since the targets are set by essentially "normalizing" history, it is best to base calibrations on a normalized environment.

Step 4: Review Model Output

The final step is to take the optimum set of calibration parameters and use them to generate a 500-scenario CAP:Link projection. In this projection we have started with initial conditions so that we can evaluate the effect of the trends. Now we are able to fully evaluate the effect of initial conditions on the optimized parameters. These results must be fully reviewed by an experienced asset simulation expert to determine the reasonableness of the results. In the end, any calibration is only as good as the credibility of the results.

4.2 Linking Assets and Liabilities in DFA Simulations

Next, we consider a hypothetical insurance line of automobile policies. We assume that these liabilities are driven by an equally weighted combination of medical inflation and legal services inflation. Using a starting liability value of $80 million, combined with the stochastic liability growth rates, we can project future liabilities. The initial asset value of $100 million can likewise be combined with the stochastic asset growth rates to project future asset values. For our analysis, we focus on the difference between the assets and liabilities – dollar surplus. The simulation renders investment and business decisions each month over the 10 year horizon.

4.2.1 Generating the Asset/Liability Efficient Frontier

We can use an asset/liability optimizer to generate an efficient frontier. The efficient frontier tells us the combination of assets that produce portfolios with the highest expected reward for a given level of risk at the end of the multi-period horizon. In this case, we have defined reward to be ending dollar surplus and risk to be the standard deviation of dollar surplus. To generate the surplus efficient frontier requires a proper multi-period DFA system. These results show the benefits of calibrating the assets and liabilities to a common set of economic factors.

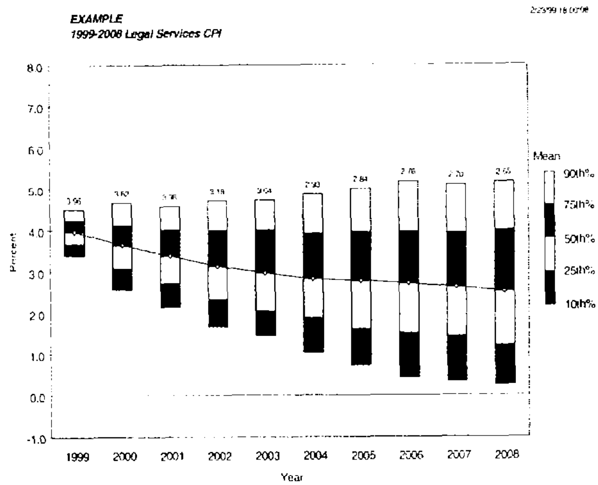

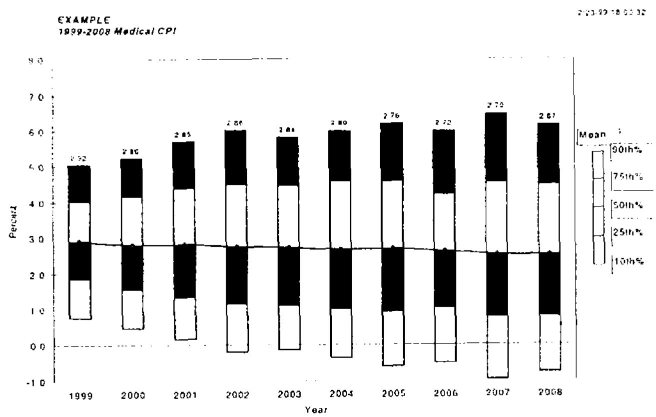

Figure 6
Projected Distributions of Medical and Legal Inflation over the Planning Horizon

4.2.2 Asset Classes and Constraints

For our analysis we have included the following asset classes and constraints. The DFA model resets the asset proportions to these value at each time step. Rebalancing the portfolio is conducted by following a fixed mix decision rule.

Asset class	Min %	Max %	Current Portfolio %
Cash	0	100	5
US Large Cap Equity	0	100	20
Bond Index	0	100	75

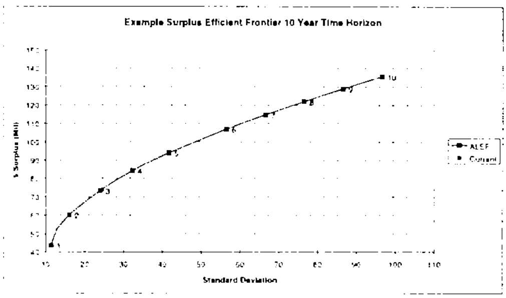

Asset Mix %	1	2	3	4	5	6	7	8	9	10	Current
Cash U.S.A	74.0	38.4	12.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	5.0
Eqty U.S.A	1.9	13.3	22.1	32.6	45.8	62.9	73.2	82.7	91.6	100.0	20.0
Bond Index	24.1	48.3	65.4	67.4	54.2	37.1	26.8	17.3	8.4	0.0	75.0
Reward	43.5	60.2	73.2	84.2	95.9	106.8	114.7	122.5	128.9	135.4	73.7
Risk	17.4	16.1	24.0	32.3	41.7	56.5	66.6	76.5	86.0	90.6	24.5

Asset/Liability Efficient Frontier (ALEF^{sm})
Figure 7

5.0 Concluding Remarks

This paper described a systematic method for calibrating a stochastic scenario generator for DFA based on an optimal fitting problem over a set of sampled scenarios. As shown by the example, the resulting calibration tool can be solved by Tabu search and related meta-heuristic approaches. Assets and liabilities should be calibrated together since there are underlying driving factors that affect the company's surplus. To properly calculate risk, we must consider both sides of the balance sheet within a DFA system. The integrated parameter estimation provides a practical method for solving this problem.

Two lines of research merit attention. The first requires the development of better ways to address the non-convex optimization model. We are currently investigating an adaptive algorithm that takes into consideration sampling errors. The goal is to solve the optimization model with the greatest degree of confidence and the least amount of sampling error. The second avenue for research is to extend the procedure to the selection of the forecasting model structure itself. Certain well-defined structural changes could possibly improve greatly our ability to generate scenarios exhibiting the desired behavior. Here, we are taking up the difficult issue of model structure error.

Notwithstanding these issues, we have shown that employing an optimization model for calibration is a practical procedure. We have illustrated the approach on a forecasting model for financial planning -- CAP:Link. We believe the approach holds promise for forecasting systems in other planning domains.

References

Berger, A. and C. Madsen (1999), "A Comprehensive System for Selecting and Evaluating DFA Model Parameters," CAS DFA Forum.

Bradford, David F. (1998) (ed.), **The Economics of Property-Casualty Insurance**, University of Chicago Press.

Broze, L., Scaillet, O. and Zakoian, J-M. (1995) "Testing for Continuous-Time Models of the Short-Term Interest Rate," *Journal of Empirical Finance*, **2**, 199-223.

Campbell, J., A. Lo, A. C. MacKinlay (1997) **The Econometrics of Financial Markets**, Princeton University Press.

Cariño, D.R. and W. T. Ziemba (1998) "Formulation of the Russell-Yasuda Kasai Financial Planning Model," *Operations Research*, July-August.

Correnti, S., S. Sonlin, and D. Isaac (1998) "Applying a DFA Model to Improve Strategic Business Decisions," *Casualty Actuarial Society Forum.*

Duffie, D. and Singleton, K. (1993) "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, **61**, 929-952.

Froot, K. and J. Stein (1998) "Risk Management, Capital Budgeting, and Capital Structure Policy for Financial Institutions: An Integrated Approach," *Journal of Financial Economics*, **47**, 55-82.

Glover, F. (1990) "Tabu Search: A Tutorial," *Interfaces*, **20**, 74-94.

Glover, F. (1995) "Tabu Search Fundamentals and Uses," Technical Report, University of Colorado, Boulder, Colorado, 80309.

Glover, F., Mulvey J.M., and Hoyland K. (1995) "Solving Dynamic Stochastic Control Problems in Finance Using Tabu Search with Variable Scaling," in *Proceedings of the Meta-Heuristics International Conference MIC-95*, Kluwer Academic Publishers, 429-448.

Gourieroux, C., Monfort, A. and Renault, E. (1993) "Indirect Inference," *Journal of Applied Econometrics*, **8**, S85-S118.

Gregory, A. and Smith, G. (1990) "Calibration as Estimation", *Econometric Reviews*, **9**, 57-89.

Hansen, L. (1982) "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, **50**, 1029-1054.

Hansen, L. and Singleton, K. (1982) "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, **50**, 1269-1286.

Holmer, M. (1998) "Integrated Asset-Liability Management: An Implementation Case Study," in **Worldwide Asset and Liability Modeling**, (eds., W.T. Ziemba and J. M. Mulvey), Cambridge University Press.

Lee, B. and Ingram, B. (1991) "Simulation Estimation of Time Series Models," *Journal of Econometrics*, **47**, 197-205.

Lowe, S.P., and J. Stanard (1996) "An Integrated Dynamic Financial Analysis and Decision Support System for a Property Catastrophe Reinsurer," *Casualty Actuarial Society Forum.*

McFadden, D. (1989) "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica*, **57**, 995-1026.

Mulvey, J. (1996) "Generating Scenarios for the Towers Perrin Investment Systems," *Interfaces*, March 1996.

Mulvey, J., R. Rush, J. Mitchell, and T. Willemain (1997) "Stratified Filtered Sampling in Stochastic Optimization," Princeton University report, SOR-97-7.

Mulvey, J., D. Rosenbaum, and B. Shetty (1996) "Parameter Estimation in Stochastic Scenario Generation Systems," Princeton University Report, SOR-96-15.

Mulvey, J. and A. Thorlacius (1998) "The Towers Perrin Global Capital Market Scenario Generation System: CAP:Link," in **World Wide Asset and Liability Modeling** by W. T. Ziemba and J. M. Mulvey (eds.), Cambridge University Press.

Mulvey, J.M., S. Correnti, and J. Lummis (1997) "Total Integrated Risk Management: Insurance Elements," Princeton University Report SOR-97-2.

Mulvey, J.M. and C.K. Madsen (1999) "Strategic Decision Making for Pension Plans Under Stress," Princeton University SOR Report.

Mulvey, J.M. and W.T. Ziemba (1998) "Asset and Liability Management Systems for long-term Investors: Discussion of the Issues," in **Worldwide Asset and Liability Modeling**, (eds., W. T. Ziemba and J. M. Mulvey), Cambridge University Press.

Pakes, A. and Pollard, D. (1989) "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, **57**, 1027-1057.