

Value-at-Risk Based Portfolio Optimization

Working Paper by
Amy v. Puelz*

* Edwin L. Cox School of Business, Southern Methodist University, Dallas, Texas 75275
apuelz@mail.cox.smu.edu, Tel: (214) 768-3151, Fax: (214) 768-4099

Value-at-Risk Based Portfolio Optimization

Abstract

The Value at Risk (VaR) metric, a widely reported and accepted measure of financial risk across industry segments and market participants, is discrete by nature measuring the probability of worst case portfolio performance. In this paper I present four model frameworks that apply VaR to *ex ante* portfolio decisions. The mean-variance model, Young's (1998) minimax model and Hiller and Eckstein's (1993) stochastic programming model are extended to incorporate VaR. A fourth model, that is new, implements stochastic programming with a return aggregation technique. Performance tests are conducted on the four models using empirical and simulated data. The new model most closely matches the discrete nature of VaR exhibiting statistically superior performance across the series of tests. Robustness tests of the four model forms provides support for the argument that VaR-based investment strategies lead to higher risk decision than those where the severity of worst case performance is also considered.

Keywords: Finance, Investment analysis, Stochastic Optimization, Value at Risk

Value-at-Risk Based Portfolio Optimization

1. Introduction

Generalized models for portfolio selection have evolved over the years from early mean-variance formats based on Markowitz's pioneering work (1959) to more recent scenario-based stochastic optimization forms (Hiller and Eckstein, 1993; Birge and Rosa, 1995; Mulvey et al, 1995; Bai et al, 1997; Vladimirov and Zenios, 1997; Cariño and Ziemba, 1998). Whether models are used for selecting portfolios for an investor's equity holdings, a firm's asset and liability combination, or a bank's derivative mix, the common thread in all model structures is the minimization of some measure of risk while simultaneously maximizing some measure of performance. In most model frameworks the risk metric is a function of the entire range of possible portfolio returns. For example, overall portfolio variance is used in a mean-variance framework while concave utility functions are applied across the set of all possible outcomes in stochastic programming frameworks. In many cases this exhaustive form of risk measurement does not accurately reflect the risk tolerance tastes of the individual or firm when there is concern with downside risk such as the risk of extremely low returns or negative cash flows.

A popularly embraced technique for measuring downside risk in a portfolio is Value at Risk (VaR). VaR is defined as the p^{th} percentile of portfolio return at the end of the planning horizon and for low p values (e.g. 1, 5 or 10) it can be thought of as identifying the "worst case" outcome of portfolio performance. Stambaugh (1996) outlines the uses of VaR as 1) providing a common language for risk, 2) allowing for more effective and consistent internal risk management, risk limit setting and evaluation, 3) providing an enterprise-wide mechanism for external regulation, and 4) providing investors with a understandable tool for risk assessment. Moreover, VaR as a practical measure of risk has been accepted by managers of firms as an

integrated and functional internal risk measure and by investors as an intuitive presentation of overall risk using a single currency valued number allowing for easy comparison among investment alternatives. The Group of Thirty, the Derivatives Product Group, the International Swaps and Derivatives Association, Bank of International Settlements, and European Union among others have all recognized VaR at some level as the standard for risk assessment. Regulators have employed VaR as a tool for controlling enterprise risk without mandates on holdings of individual instruments. For example, the 1998 Basle Capital Accord proposes a bank's required set aside capital for market risk based on internal VaR estimates and the National Association of Insurance Commissioners (NAIC) also requires the reporting of VaR.

The academic and practitioner literature is now filled with tips, techniques, and problems associated with the derivation and validity of a portfolio's VaR *ex post*. One interesting perspective of VaR is how it may be applied to *ex ante* portfolio allocation decisions. If firms are to be regulated and performance is to be evaluated based on VaR, then it should be considered when developing a strategy for project or investment selection. Further, if firms do make decisions in a VaR context, what are the broader implications for a firm's risk exposure?

The literature has been relatively silent on the topic of optimization model development under VaR because of its discrete nature which is difficult to incorporate in traditional stochastic model forms. Uryasev and Rockafellar (1999) propose a scenario-based model for portfolio optimization using Conditional Value at Risk (CVaR) which is defined as expected value of losses exceeding VaR. Their optimization model minimizes CVaR while calculating VaR and in the case of normally distributed portfolio returns; the minimum-CVaR portfolio is equivalent to the minimum-VaR portfolio. Kalin and Zagst (1999) show how VaR can be derived from

models with volatility or shortfall constraints. Neither of these papers, however, proposes a model that incorporates VaR as the risk metric for portfolio evaluation.

Relative to how the incorporation of VaR in investment decisions influences the firm's risk taking behavior, Basak and Shapiro (1999) show theoretically that optimal decisions based on VaR result in higher risk exposure than when decisions are based on expected losses. Specifically, they show that when losses occur they will be greater under a VaR risk management strategy. They suggest that this is a significant shortcoming of VaR-based policies and propose that alternative measures, based on the expected value of losses, be incorporated in risk management strategies.

The purposes of this paper are twofold: 1) the development of a model form that derives optimal *ex ante* VaR portfolios and 2) the presentation of empirical and simulated evidence that VaR-based strategies are riskier than those where expected losses are incorporated in the risk metric. Three of the models developed for VaR portfolio selection are direct extensions of models developed primarily for use with other risk metrics and incorporate, in some form, expected losses in addition to the probability of loss as in VaR. The first of these three is based on the classic mean-variance (MV) framework, the second is the minimax (MM) model developed by Young (1998) and the third is derived from Hiller and Eckstein's (1993) scenario-based stochastic programming (SP) model for asset-liability management. The fourth model form is new and combines stochastic programming with an aggregation/convergence technique (SP-A) developed by Puelz (1999).

Portfolios are generated that meet VaR specified standards for risk (henceforth referred to as VaR-feasible portfolios) and optimal portfolios are compared across model forms. Overall the best model for incorporating a VaR risk strategy is the SP-A model which generates VaR-

feasible portfolios with statistically significant higher returns across all empirical and simulated tests. Robustness tests of resulting portfolio performance reveal that employing a VaR strategy leads to higher levels of risk taking than strategies where the severity of loss is embedded in risk assessment.

The next section summarizes much of the literature addressing VaR calculation and validity, and provides background necessary to the model test methodology, results, and performance. The four model forms to be tested are presented in section three and the results of empirical and simulated performance tests are reported in section four. Concluding remarks are put forth in section five.

2. VaR - Measurement and Validity

VaR was first widely accepted as a risk measurement standard at derivative trading desks where it served as a technique to combine and measure disparate risks. Since that time it has become the standard at other trading operations, financial institutions, insurers, and firm cash-flow management operations (Simons, 1996; Beder, 1995). There are three basic techniques employed to measure VaR: mean-variance, historical simulation and Monte Carlo simulation.

The first technique for VaR estimation involves a closed-form solution using the variance/covariance matrix of security returns. This approach is the most simplistic and in many ways the easiest to understand as it relates directly back to mean-variance efficient portfolio derivation. However, simplicity, as would be expected, comes at a cost. The most significant limitation of the MV approach is an implied assumption in most models of return normality. This assumption has been proven invalid as most securities exhibit return behavior characterized by high levels of kurtosis resulting in "fat tails" or leptokurtic distributions. Securities with

leptokurtic return distributions are more likely to realize unusual returns (e.g. +/- 4 standard deviations) than comparable normally distributed return securities. Duffie and Pan (1997) show the impact nonlinear returns or "jumps" and stochastic volatility have on increasing return kurtosis. Lucas and Klaassen (1998) prove that portfolio allocations will be overly risky or overly prudent under an assumption of return normality when returns are actually leptokurtic.

The MV approach proposed for VaR *ex ante* portfolio selection in this paper involves searching the efficient frontier until the desired VaR is achieved. It is possible that if convergence occurs at the minimum-variance or maximum-return portfolio, the MV model derived solution will be overly risky or overly conservative, respectively. If convergence occurs at an efficient portfolio between the two extremes, target VaR will be met. In theory, if this occurs and returns are symmetrically distributed, the MV optimal portfolio will be the optimal maximum-return, VaR-feasible portfolio. In practice, however, given imperfect return distribution specification and MV model sensitivity to small changes in input parameters, expected return on optimal MV portfolios will be below the maximum attainable level under a VaR metric.

The second VaR measurement technique uses historical returns under a simulation approach. Intra-day, daily, weekly, monthly or quarterly historical return data are used as scenarios or possible future realizations to directly calculate the portfolio's VaR. The benefit of this approach over a mean-variance approach is that no assumptions about the distribution of security returns are required. It is assumed, however, that historical trends will continue into the future resulting in models that are less sensitive to changing market conditions. Sample size is another issue when using historical simulation. Larger scenario sample sets must be drawn from increasingly older return histories compounding the problem of model insensitivity to changing

markets. In the historical data driven tests to follow the impact of changing market conditions on model robustness is evident.

The third technique for measuring VaR, Monte Carlo simulation, allows for complete flexibility with regard to security return distributions. In addition, the scenario sample set is not limited by historical realizations as securities are priced based on algorithms and/or heuristics imbedded in the model. Monte Carlo frameworks have been shown to provide the best estimates for VaR (Pritsker, 1997; Lucas and Klaassen, 1998). At the same time these models are by far the most computationally intensive due both to the internal pricing mechanisms and to the large scenario sets necessary to adequately represent highly unlikely market swings. A single index Monte Carlo simulation model is used in this paper allowing for model performance evaluation given known return distribution characteristics.

3. VaR Based Model Forms

The four models to follow, all variations of general-form portfolio selection models, are employed across a set of scenarios representing security returns. Each scenario represents one possible joint realization of uncertain returns over the planning horizon. In the MV model, scenarios are used to derive the input variance/covariance matrix, and in the other models the scenarios are input directly. Each of the models is developed under the assumptions of a single-period planning horizon and no short sales. These assumptions can be relaxed in all but the mean-variance model and are in place to allow for comparisons across model forms. The target "worst case" return for VaR is represented in all models as a p^{th} percentile return of R^* .

3.1 MV Model

The MV model is defined as follows

$$\min \quad \sum_i \sum_j x_i x_j \text{COV}_{ij} \quad (1a)$$

$$\text{ST} \quad \bar{C} X^T \geq G \quad (1b)$$

$$P X^T \leq 1 \quad (1c)$$

$$X \geq 0, \quad (1d)$$

where X is the $1 \times N$ vector of security allocation proportions $\{x_j\}$, COV_{ij} is the covariance between securities i and j , \bar{C} is the $1 \times N$ vector of security expected returns $\{\bar{c}_j\}$, P is the $1 \times N$ vector of security prices $\{p_j\}$, and G is the minimum required expected return. To implement the MV model in a VaR decision framework, G is varied systematically from the expected return of the minimum-variance portfolio to that of the maximum-return portfolio until the target VaR is reached.

There is no guarantee that the MV approach will converge at a VaR-feasible solution. It is possible that the MV approach will converge at the minimum-variance portfolio with the p^{th} percentile return less than the target R^* . The higher the level of risk aversion the greater the probability this will happen and portfolio risk will be greater than that specified by the target VaR. Likewise, it is also possible the MV approach will converge at the maximum-return portfolio with the p^{th} percentile return greater than the target R^* resulting in an overly prudent portfolio relative to the desired VaR. The probability of this occurring increases as the level of risk aversion decreases. In order for the MV approach to derive a portfolio that is neither overly risky nor overly prudent, R^* must lie somewhere between the p^{th} percentile of the maximum-return portfolio and the p^{th} percentile of the minimum-variance portfolio.

It is important to note that, given symmetric portfolio return distributions, the MV approach will in theory derive the maximum expected-return portfolio for a target VaR if R^* lies in the allowable range described above. However, in practice, given a limited scenario sample set, the MV model will become less reliable as the level of kurtosis increases. The MV approach has been shown to be very sensitive to input parameter misspecification (Koskosidis and Duarte, 1997). If sample scenario returns are skewed because of a few extremes at either end, the MV solution is more likely to converge at the minimum-variance portfolio if a negative skew exists and to converge at an overly prudent portfolio if a positive skew exists. In the tests to follow, R^* is set at a high value to insure the MV model will not converge at the minimum-variance portfolio. However, the MV model because of leptokurtic return distributions will consistently converge at overly prudent portfolios relative to VaR-feasible maximum-return portfolios.

3.2 MM Model

The MM portfolio framework proposed by Young (1998) is of the form

$$\max \quad m \tag{2a}$$

$$\text{ST} \quad C(\xi)X^T - m > 0 \quad \forall \xi=1 \text{ to } S \tag{2b}$$

$$\bar{C}X^T \geq G \tag{2c}$$

$$PX^T \leq 1 \tag{2d}$$

$$X \geq 0, \tag{2e}$$

where m is the return on the minimum return portfolio, and $C(\xi)$ is the $1 \times N$ vector of security returns under scenario ξ , $\{c_j(\xi)\}$. Young proposes this model as a practical approach to portfolio selection that financial managers can implement given standard computing technology. The

linear programming structure also allows for the incorporation of special form side constraints. In extensive testing of the MM model, its performance is found to be similar to that realized using a MV model framework. For these reasons and because of the discrete nature of the risk metric incorporated in the model, the MM model is a candidate for implementation under a VaR metric. This is accomplished in the same manner as in the MV framework by adjusting G until a target VaR return of R^* is attained.

There are two potential drawbacks to the implementation of the MM model under a VaR framework. First, model size increases directly with the number of scenarios. Second, at the optimal solution point in the MM model only the worst-case scenario constraint in (2b) is binding. Ideally the lower p^{th} percentile of portfolio constraints would be binding at the optimal solution point. This implies that the risk profiles for optimal MM portfolios are overly dependent on the worst case scenario rather than on the entire set of worst case scenarios below the p^{th} percentile. In the case of highly leptokurtic distributions (or for that matter negatively skewed distributions) the MM approach will derive overly prudent portfolios relative to VaR specified risk.

3.3 SP Model

Stochastic programming frameworks have been found to be quite robust for the class of portfolio selection problems (Cariño et al, 1994; Golub et al, 1995; Holmer and Zenios, 1995; Koskosidis and Duarte, 1997). The SP approach proposed by Hiller and Eckstein (1993) provides the most adaptable framework for VaR *ex ante* portfolio allocation. The model is of the form

$$\max \quad \bar{C}X^T - \lambda \sum_{\xi} \pi(\xi) d(\xi, X) \quad (3a)$$

$$ST \quad C(\xi)X^T + d(\xi, X) \geq R^* \quad \forall \xi = 1 \text{ to } S \quad (3b)$$

$$PX^T \leq 1 \quad (3c)$$

$$X \geq 0, \quad (3d)$$

where $d(\xi, X)$ is the amount below a return of R^* for each scenario ξ . The expected amount below R^* is represented in the objective function as $\pi(\xi)d(\xi, X)$, where $\pi(\xi)$ is the probability of scenario ξ . This risk term which is the CVaR for the portfolio (Uryasev and Rockafellar, 1999) is weighted by λ , where a higher λ corresponds to a higher level of risk aversion. Hiller and Eckstein propose this model as a stochastic dedication model for fixed-income portfolios. In their original model for asset/liability management, R^* is set at zero creating a risk term that reflects the probability and severity of insolvency across all scenarios. Before Hiller and Eckstein's model the implementation of such a lower partial mean risk metric was considered computationally intractable in stochastic programming. The particular formulation proposed by Hiller and Eckstein is easily adapted to a VaR framework by setting R^* to the target worst case return and varying the risk parameter λ until $\sum_{\xi} \pi(\xi)$ is equal to $p/100$, the desired probability density function at R^* .

Several problems can be encountered, however, in implementing SP for portfolio selection under VaR. First, SP models are computationally intensive and in many cases become intractable given a realistic set of assumptions, a reasonable number of securities, and a multi-time period planning horizon (Dahl et al, 1993). Like the MM model framework, SP model size increases directly with the number of scenarios becoming a significant implementation issue in a VaR framework if the probabilities of unlikely and possibly catastrophic events are to be represented in the scenario set. By not incorporating these risks, models will fail to control for

the types of risks that served as a major impetus for the acceptance of VaR (McKay and Keefer, 1996). Second, the SP model will converge at overly prudent portfolio relative to specified VaR, when portfolio returns are highly leptokurtic. The impact on return is not as significant as seen in the MM model, but because worst case return scenarios are all included in the objective function, to achieve the desired VaR for average worst case return, expected return performance is sacrificed. This means that the SP framework, like the MM and MV approaches, embed expected losses in the measurement of “worst case” performance.

3.4 SP-A Model

The new model form, SP-A, addresses the problem of model size given large scenario sets and the impact the level of kurtosis and return distribution misspecification have on overly prudent portfolios under VaR. The SP-A technique provides a tractable and robust modeling approach that in simulated and empirical tests derives superior VaR-feasible portfolios compared to the other models. Robustness tests on these VaR optimal portfolios illustrate that VaR-based strategies for investment selection lead to riskier investment plans than when the severity of expected losses is considered as in the MM, MV, and SP models.

The SP-A selection technique uses Hiller and Eckstein's basic model form and implements it using scenario aggregation in combination with a convergence algorithm to derive VaR-feasible portfolios. The aggregation/convergence technique proposed as a tractable alternative for solving stochastic programming models was first applied to asset-liability matching under quadratic utility (Puelz, 1999). In the case of VaR-based optimization, aggregation not only reduces model size making it feasible to solve without decomposition but

also allows for the incorporation of VaR in its pure discrete form. The SP-A model under VaR is of the form

$$\max \quad \bar{C}X^T - \lambda \sum_k \pi_k d_k \quad (4a)$$

$$\text{ST} \quad C_k X^T + d_k \geq R^* \quad \forall k = 1 \text{ to } K \quad (4b)$$

$$PX^T \leq 1 \quad (4c)$$

$$X \geq 0, \quad (4d)$$

where C_k is the $1 \times N$ vector of aggregated security returns for return category k derived as

$$C_k = E_{\xi} [C(\xi)] \quad \forall \xi \in \text{Category } k, \quad (4e)$$

d_k is the amount below a return of R^* for category k , and π_k is defined as

$$\pi_k = \sum_{\xi} \pi(\xi) \quad \forall \xi \in \text{Category } k. \quad (4f)$$

In the SP-A approach, the optimal portfolio is derived through an iterative procedure illustrated in Figure 1.

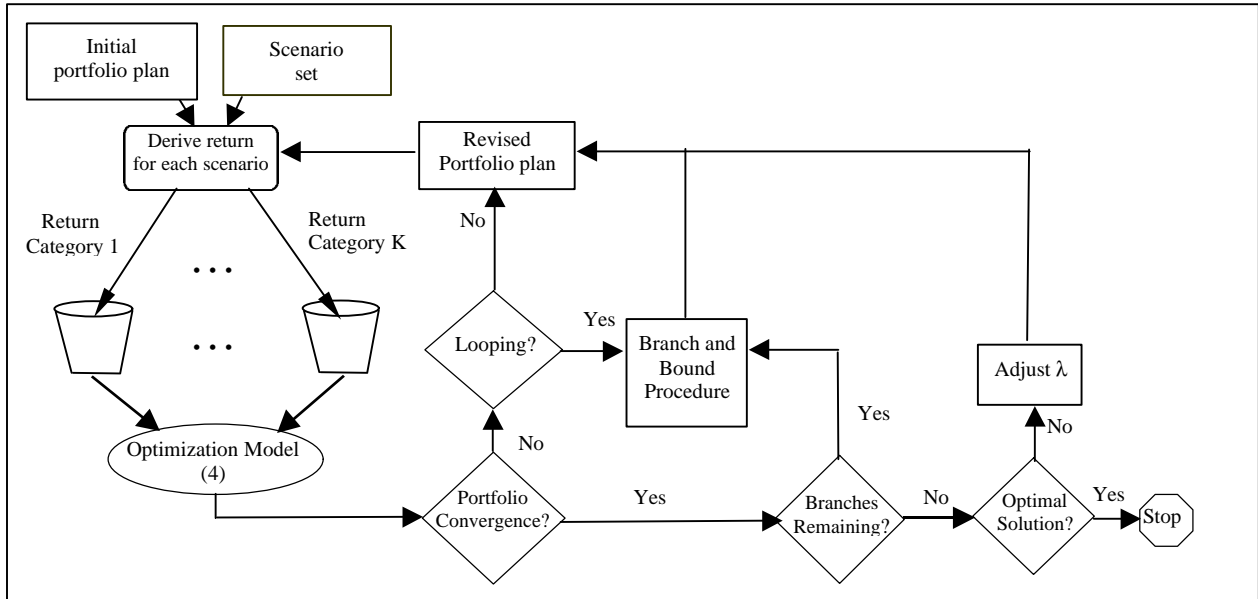


Figure 1: SP-A process for VaR-based portfolio allocation

First, the returns of an initial portfolio plan are derived across all scenarios. These returns are used to aggregate the scenarios into categories. The SP-A model in (4) is solved to derive a revised portfolio plan and the revised plan is used to re-aggregate scenarios and (4) solved again. This process of re-aggregating and revising portfolios continues until the model converges to the optimal plan. Convergence occurs when (4) produces the same portfolio plan as that used to aggregate scenarios. If the model does not converge, then all previous plans are compared to determine if any are identical indicating that the model is looping between a set of plans and will not converge. If true, a modified branch and bound procedure is implemented to search the solution space deriving the best, VaR-feasible converged solution. The final converged solution is evaluated relative to $\sum_{\xi} \pi(\xi)$, as with the SP model, adjusting λ until the maximum-return, VaR-feasible portfolio is obtained. The modified branch and bound procedure and the adjustment process for λ are discussed in detail in the appendix.

An important component of the SP-A model is determining the ranges for the return aggregation categories. With equity securities exhibiting continuous return distributions the ranges will depend on the target VaR with most falling below or slightly above R^* . All high level returns (i.e. above the 50th percentile) would be aggregated into a single large category. This approach allows for risk assessment focused on the worst or below-average outcomes, which is especially well suited to a VaR metric. The aggregation or categorization process based on scenario returns is reflective of the VaR measure itself, which is by nature a categorical worst case measure rather than a continuous one. Unlike the MM or SP models, the SP-A incorporates average returns in the objective function making it less sensitive to extremely low returns when returns distributions are highly leptokurtic.

4. Model Tests

This section presents the results of various tests of the four models for portfolio selection. These tests are structured to compare the models given an R^* such that the MV model converges at an optimum other than the minimum-variance or maximum-return portfolio. Empirical tests use an historical simulation approach to derive portfolios of international equity indices. These tests over a fifteen-year period illustrate the relative model performance and the impact of changing market conditions on model robustness. A Monte Carlo approach using simulated returns based on a single index model is used to test the models given varying levels of return kurtosis. A discussion of the both empirical and simulated tests is followed by a summary of the computational experience for the SP-A model.

4.1 Empirical Tests

The first set of tests uses monthly returns for thirteen national indices for the period from January, 1984 through February, 1999. (Data were collected from Datastream's global equities databases.) All returns were converted to US dollar equivalent returns based on historical exchange rates. Summary statistics for the empirical data are found in Table 1.

Table 1: Descriptive statistics for national indices - January 1984 to February 1999

	Australia	Canada	Germany	Hong Kong	Japan	Singapore
Mean	0.79%	0.55%	1.23%	1.61%	0.90%	0.59%
Std. Dev.	6.79%	4.68%	5.71%	8.82%	7.59%	7.24%
Kurtosis	6.37	4.44	1.34	5.19	0.79	4.26
Skewness	-1.05	-0.72	-0.09	-0.62	0.35	-0.15

	S. Africa	UK	France	Italy	Belgium	Denmark	US
Mean	0.40%	1.24%	1.45%	1.45%	1.48%	1.06%	1.18%
Std. Dev.	8.21%	5.42%	6.15%	7.61%	5.45%	5.43%	4.13%
Kurtosis	2.26	1.39	0.51	0.45	3.92	0.92	5.24
Skewness	-0.12	-0.18	-0.03	0.45	0.54	0.16	-0.94

For each trial, fifty consecutive months of returns were used as the input scenario set to run each of the four models. The resulting portfolio performances were examined using the next consecutive twelve months of returns. Results are presented across both the fifty months used to run the models (in-sample scenarios) and the hold-out sample of twelve months (out-sample scenarios). Eleven trials were executed with the first trial beginning in January 1984 and the last trial beginning in January of 1994. The beginning date of each trial was incremented by twelve months from the previous trial.

All models were run with a target percentile for VaR, p , of 10 and R^* set to the average of the 10th percentile of returns across in-sample scenarios for the minimum-variance and maximum-return mean-variance efficient portfolios. This assured the convergence of the MV model at or close to the target VaR rather than at the upper or lower bound for efficient frontier risk. The 10th percentile for VaR was selected to facilitate testing using smaller sample scenario sets because the smaller the percentile used, the larger the sample scenario set required to capture unusual outcomes. In terms of relative comparisons of model performance, the conclusions drawn are applicable to any target VaR percentile.

The SP-A model initial portfolio for aggregation was set at the MV optimal portfolio and model convergence was assumed when security allocations of the revised portfolio varied by no more than 2.5% from the portfolio used for aggregation. Two aggregation categories with equal return ranges were below R^* and three categories with equal ranges were between R^* and the 50th percentile of portfolio returns.

The summary statistics for the eleven-period test horizon are presented in Table 2. The superscripted numbers in Table 2 are the p-values for two-tailed t-tests of a hypothesized mean

difference of zero between each statistic for the MV, MM and SP models and the comparable statistic for the SP-A model. The last statistic, average VaR attainment, is the average proportion R^* exceeds the p^{th} percentile of portfolio returns across the eleven trials. A negative number for VaR attainment indicates the amount VaR-specified risk was exceeded in the optimal portfolio. A VaR-feasible portfolio will have a nonnegative number for VaR attainment.

Table 2: Summary statistic for empirical tests

Across in-sample scenarios	MV	MM	SP	SP-A
Average return	2.13% ^{0.18}	1.86% ^{0.03}	2.08% ^{0.01}	2.21%
Average standard deviation	5.30% ^{0.11}	5.30% ^{0.23}	5.31% ^{0.00}	5.57%
Average VaR attainment (p^{th} percentile - R^*)/ R^*	0.01 ^{0.21}	-0.15 ^{0.05}	0.02 ^{0.65}	0.02
Across out-sample scenarios	MV	MM	SP	SP-A
Average return	0.70% ^{0.34}	0.86% ^{0.38}	0.76% ^{0.54}	0.62%
Average standard deviation	6.05% ^{0.08}	5.90% ^{0.20}	5.90% ^{0.00}	6.34%
Average VaR attainment (p^{th} percentile - R^*)/ R^*	-0.67 ^{0.12}	-0.60 ^{0.19}	-0.64 ^{0.08}	-0.78

Examination of performance across in-sample scenarios allows for comparison of the models given perfect information about future scenarios. The MV and SP-A models outperform the MM and SP models with significantly higher expected returns at the same level of risk relative to VaR. The MM model and to a lesser extent the SP model overcompensate for VaR when returns exhibit excess levels of kurtosis as the indices used in the test. The MV model average return was lower than that for SP-A but not at a statistically significant level. All the models except the MM model derived VaR-feasible solutions across in-sample scenarios.

A comparison of average standard deviation values across in-sample scenarios reveals that the SP-A portfolios have higher overall variability of returns compared to other models. The MM, and SP models derive more prudent, lower variability portfolios because of the incorporation of worst case return values directly in the objective function. In the case of the MV model overall variance is used to define the efficient frontier. The aggregation of returns in SP-A reduces the indirect incorporation of a secondary continuous risk metric based on overall variability or expected value of losses and derives portfolios with higher expected returns given the same target VaR specifications.

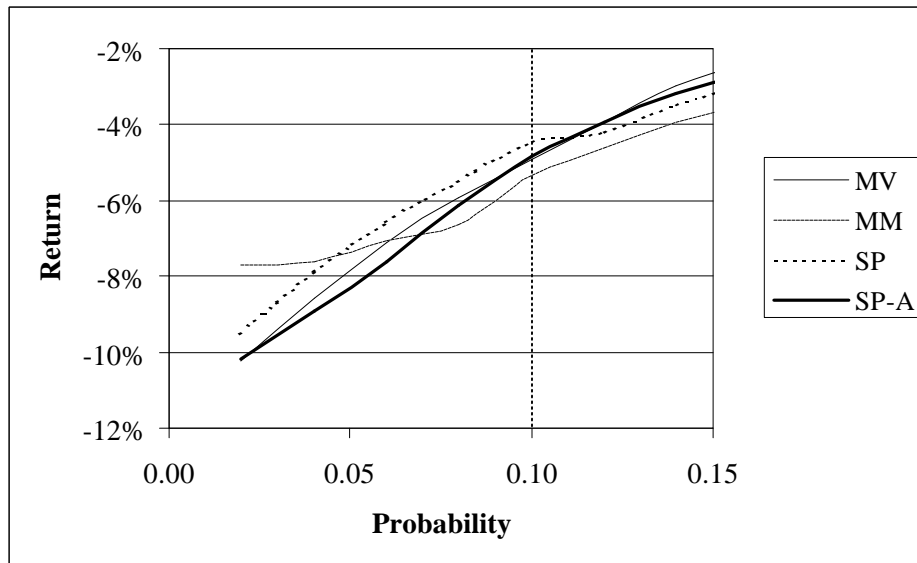


Figure 2: Cumulative probability of portfolio returns across in-sample scenarios in empirical tests.

Figure 2 plots the lower tail of the cumulative return distribution for in-sample scenarios averaged across the eleven trials. The MM model, where minimum return is controlled, clearly results in the best of the “worst case” return outcomes. SP, and to a lesser extent MV, both have higher returns for percentiles below the 10th due to the nature of their objective functions. This plot reveals that under a VaR strategy, where only the pth percentile loss is controlled, larger

losses below the p^{th} percentile are realized than under a strategy where the severity of losses below the p^{th} percentile are controlled. This outcome supports the theoretically based conclusions drawn by Basak and Shapiro (1999).

Upon examination of the performance across out-sample scenarios in Table 2, it is clear that portfolio performance is not robust to the holdout sample set of future scenarios. In this case the behavior of future returns is not well described by observed behavior dating back 50 months. However, using a smaller sample set of scenarios would significantly limit the representation of worst case returns. It is difficult to draw a large enough sample from historical returns that is still predictive of future return behavior. For this reason historical simulation for VaR-based portfolio selection is limited by the inability to generate portfolios that will perform as expected into the future.

It is apparent, however, that the SP-A model expected returns are lower (although not highly statistically significantly lower) than the other models. In this case, the more prudent strategies derived with consideration of expected losses, as with MV, MM and SP, fare better when poor information is available relative to future return behavior. The next set of simulated tests control for the quality of information captured in model scenarios by using a Monte Carlo simulation technique where known return distribution characteristics are used to generate both in-sample and out-sample scenarios sets.

4.2 Simulated Tests

In simulated tests, scenario returns for ten different securities were generated using a single index model given as

$$c_j(\xi) = a_j + b_j m(\xi) + s_j e_j(\xi), \quad (5)$$

where $m(\xi)$ is a student-t distributed random variable with a mean of 0 and v degrees of freedom, and $e_j(\xi)$ is a normally distributed random variable with a mean of 0 and standard deviation of 1. Both b_j and s_j ranged from 0.1 to 5.0 and were independently and randomly assigned to each of the ten securities. The constant a_j ranged from 1 to 3 with the lower values assigned to securities with lower variability across the scenario set and higher values assigned to those with higher variability.

The simulated tests focus on the impact the level of kurtosis has on model performance by varying the student-t distribution degrees of freedom, v , in the return generation process. Each simulation run involves creating 200 scenarios, 100 of which are used to run each of the four models (in-sample scenarios). The remaining 100 scenarios are used to examine the robustness of model performance (out-sample scenarios). The target VaR percentile is set at 10 and R^* is derived in the same manner as described in the empirical test section. Fifty trials are executed for each simulation run with identical parameters and ranges for return categorization. The results of two sample simulation runs are presented to summarize the findings of these tests.

In the first simulation run, v is set to 3 creating security return distributions with relatively high average kurtosis of 8.0. A summary of the simulation statistics across the fifty trials is given in Table 3. The superscripted values in Table 3, as in Table 2, are the p-values for two-tailed t-tests of a hypothesized mean difference of zero between each statistic for the MV, MM and SP models and the comparable statistic for the SP-A model. For both in-sample and out-sample scenario sets, the SP-A model outperformed the alternative models with statistically significantly higher expected returns. Evidenced by relative portfolio return standard deviation values, higher returns for the SP-A model were realized by controlling only for VaR rather than indirectly controlling for overall variability or expected value of losses, as in the other models.

The increased risk realized by using a VaR strategy is seen in the VaR attainment figures across out-sample scenarios. Attainment of target VaR is lower for SP-A because it incorporates an isolated measure of risk focused on the p^{th} percentile of losses rather than the severity of losses below the p^{th} percentile. The SP-A model derives a more aggressive portfolio that will be less sensitive to potential loss outcomes below the p^{th} percentile of losses.

Table 3: Summary statistic for simulated tests with high return kurtosis ($v = 3$)

Across in-sample scenarios	MV	MM	SP	SP-A
Average return	2.64% ^{0.00}	2.33% ^{0.00}	2.66% ^{0.00}	2.71%
Average standard deviation	4.87% ^{0.00}	4.34% ^{0.00}	5.00% ^{0.00}	5.18%
Average VaR attainment (p^{th} percentile - R^*)/ R^*	0.00 ^{0.00}	0.00 ^{0.01}	0.01 ^{0.68}	0.02
Across out-sample scenarios	MV	MM	SP	SP-A
Average return	2.21% ^{0.00}	2.03% ^{0.00}	2.24% ^{0.00}	2.29%
Average standard deviation	5.31% ^{0.00}	4.64% ^{0.00}	5.45% ^{0.00}	5.66%
Average VaR attainment (p^{th} percentile - R^*)/ R^*	-0.55 ^{0.00}	-0.10 ^{0.00}	-0.65 ^{0.01}	-0.77

In Figure 3 the plot of the lower tail of the cumulative return distribution for in-sample and out-sample scenarios averaged across the 50 trials are shown. Percentiles below the 10th are lowest for SP-A as seen in the empirical tests; again providing support for the argument that VaR-based strategies are more likely to result in higher losses when losses occur.

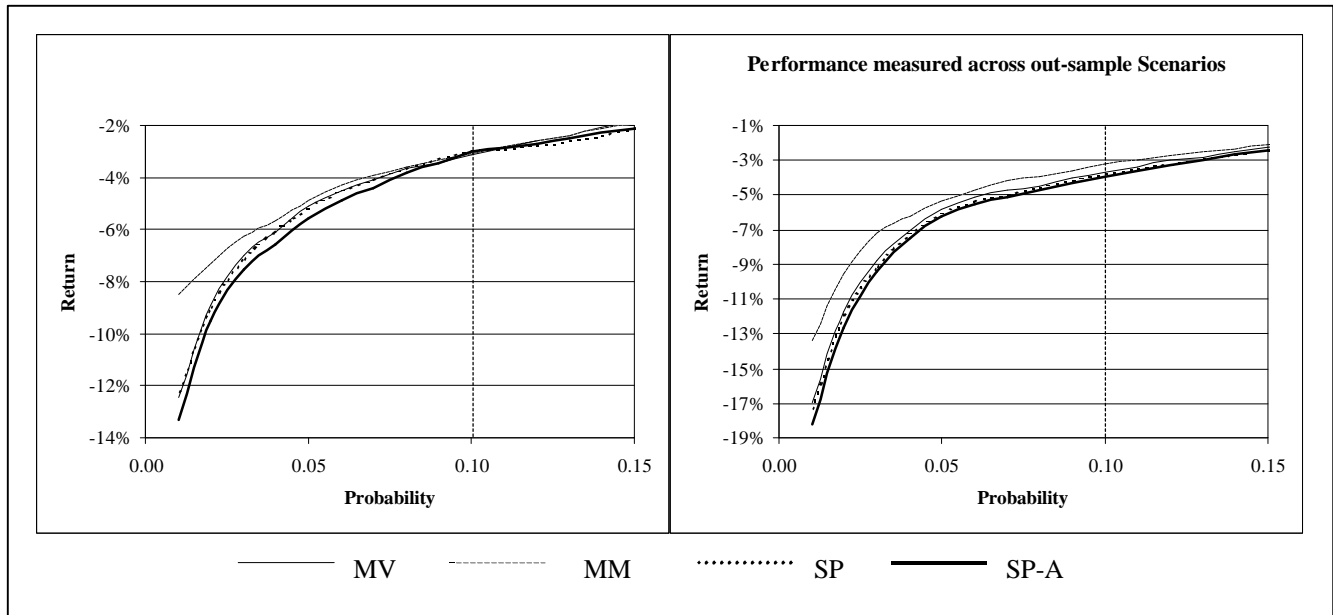


Figure 3: Cumulative probability of average portfolio returns for in-sample and out-sample scenarios across 50 simulated trials with high-level return kurtosis ($v=3$).

In the next set of simulation runs v is increased to 10 creating return distributions with relatively low average kurtosis of 3.4. The summary statistics across the fifty trials are presented in Table 4 with the superscripted values again representing p-values for a hypothesized mean difference of zero. Given lower kurtosis, the SP-A model still significantly outperforms the other models for the set of in-sample scenarios. The MM, MV and SP models average expected returns are significantly below those for the SP-A model but percentage differences are not as large as those seen in the case of high kurtosis presented in Table 3. For example, the MM model expected return is 16% below the SP-A model across in-sample scenarios with high kurtosis ($v = 3$) and only 8% below with low kurtosis ($v = 10$). This is because with lower levels of kurtosis fewer extreme returns are present in the scenario sample set to influence the selection in the MM, MV, and SP models. The SP-A still has higher portfolio standard deviation and

lower VaR attainment than other model forms but again the percentage difference is less with lower kurtosis.

Table 4: Summary statistic simulated tests with low return kurtosis ($v = 10$)

Across in-sample scenarios	MV	MM	SP	SP-A
Average return	2.95% ^{0.00}	2.80% ^{0.00}	2.97% ^{0.00}	3.03%
Average standard deviation	4.15% ^{0.00}	3.88% ^{0.00}	4.25% ^{0.00}	4.42%
Average VaR attainment (p^{th} percentile - R^*)/ R^*	0.00 ^{0.00}	0.00 ^{0.00}	0.01 ^{0.02}	0.02
Across out-sample scenarios	MV	MM	SP	SP-A
Average return	2.59% ^{0.00}	2.51% ^{0.00}	2.63% ^{0.15}	2.65%
Average standard deviation	4.38% ^{0.00}	4.08% ^{0.00}	4.45% ^{0.00}	4.63%
Average VaR attainment (p^{th} percentile - R^*)/ R^*	-0.73 ^{0.00}	-0.29 ^{0.00}	-0.70 ^{0.00}	-1.01

4.3 Computational Experience - SP-A Model

The SP-A model under VaR involves running the model until convergence occurs and adjusting the risk parameter, λ , until the optimal VaR-feasible portfolio is derived. The average number of adjustments of λ before the optimal SP-A portfolio was derived was 51 in the empirical tests and 38 in the simulated tests. The average number of times the SP-A optimization model (4) was executed for each adjustment of λ of was 161 for the empirical tests and 552 for the simulated tests. The average number of branches that were added to (4) in the modified branch and bound procedure for model convergence was 8 for the empirical tests and 11 for the simulated tests.

Overall SP-A is more computational intensive given the size of the sample scenario sets used in these tests. However, as the size of the scenario set increases, the size of the SP-A model in (4) will remain constant for a given number of return categories while the SP model in (3) and MM model in (2) will increase. Under a VaR framework the number of scenarios necessary to adequately represent the critical lower percentile of the return distribution will increase as the VaR target percentile decreases (level of risk aversion increases) and as the level of kurtosis increases. In addition, as VaR models are extended from a single-period to a multi-period planning horizon, and as the number of candidate securities to be evaluated increases, the number of scenarios required to adequately represent the joint distributional characteristics of returns significantly increases (Dahl et al 1993). Given realistically sized, large-scale problems, the MM and SP model frameworks require extensive computing resources and possibly a structure suitable for model decomposition for solution generation. This is not the case with the SP-A framework.

5. Conclusions

VaR is a widely reported and accepted measure of risk across industry segments and market participants. The models proposed in this paper are designed to provide decision frameworks for evaluating investment decisions given a VaR metric. If perfect information is available about future return behavior the SP-A model clearly outperforms the other models as demonstrated by VaR-feasible portfolios with significantly better performance across all in-sample scenario sets in simulated and empirical tests. In comparing performance across out-sample scenarios sets, the SP-A model performs better or on par to the other models relative to expected returns but is less robust in meeting target VaR specifications. This can be attributed to

the more aggressive portfolios with higher overall return variability derived with the SP-A model given the same VaR. In general, VaR optimal portfolios are more likely to incur large losses when losses occur. The more prudent approaches that incorporate the expected value of losses are more likely to not exceed VaR risk specifications when examined relative to holdout scenarios. The MM and MV models may not, however, converge at a VaR-feasible solution resulting in overly prudent or risky portfolios relative to target VaR specifications

The percentage increase in returns realized using a VaR strategy relative to more prudent strategies that incorporate the expected value of losses is positively correlated with the level of kurtosis in returns. Likewise the impact of a VaR strategy on excessive exposure to large losses is also positively correlated with the level of return kurtosis. In general, a VaR strategy becomes more risky the lower the level of return distribution specification in the model scenarios which can be caused by highly leptokurtic returns or the use of historical returns.

Implementation issues become critical as the degree of return distribution kurtosis, level of risk aversion, periods in the planning horizon or securities in the selection set increases. In any of these cases the number of scenarios necessary to adequately specify joint return distribution behavior increases. The sizes of the SP and MM models increase directly with the number of scenarios and can become intractable for large-scale implementations. By contrast, the SP-A does not suffer from this drawback, since model size is not impacted by the scenario count, which is encouraging for practical implementations.

6. Appendix - Solution procedures for the SP-A model

The solution procedure for the VaR-based SP-A model involves implementing a branch and bound procedure when looping occurs, and adjusting λ and its bounds until the solution space has been searched.

6.1 Branch and Bound Procedure

Looping occurs in the SP-A model when a portfolio plan reappears resulting in cycling between a set of portfolio plans. When this occurs the following steps are taken:

- 1a) The variable for the security with the maximum difference in allocation percentage for the two most recently derived portfolios in the loop is selected as the branching variable.
- 2a) A less-than constraint is added to the optimization model in (4) with the LHS equal to the selected branching variable and the RHS equal to the minimum value that holds the branching variable constant from run to run of the SP-A model.

Branches are added to (4) in this manner until the SP-A model converges. If the converged portfolio's expected return is greater than the expected return of the current best solution and the portfolio's p th percentile is not less than the target R^* (a VaR-feasible portfolio) then the solution replaces the current best solution.

Once the model converges, if branching constraints remain in (4) they are changed and/or removed as follows:

- 1b) The most recent less-than constraint is terminated and replaced by an identical greater-than constraint.
- 2b) All constraints subsequent to the constraint changed in 1b) are removed.

Important to the tractability of the branch and bound procedure is branch termination based on the current best solution. Branches are terminated if the maximum expected return for portfolios in a loop, \bar{CX}^T , is less than the current best portfolio expected return. In this case, λ is too high on this branch to generate a solution better than the current best.

6.2 Adjusting λ and Bounds

When all branches are removed, λ is adjusted if the final optimal solution has not been found. The optimal solution has been found if the upper and lower bounds for λ , λ_U and λ_L respectively, are equal. If this is not the case, λ and its bounds are adjusted according to one of the following steps:

- 1d) If a new best solution was found since the last adjustment of λ , then λ_U is set to the current λ and the new λ is set midway between λ_L and new λ_U . In this case a VaR-feasible solution has been found and searching for higher values of λ (higher levels of risk aversion) will not uncover a superior solution.
- 2d) If a new best solution is not found, and λ is greater than λ_L , then λ is adjusted down by a factor of $(\lambda - \lambda_L) / 2$.
- 3d) If λ is equal to λ_L , λ_L is adjusted upwards by a factor of $(\lambda_U - \lambda_L) / 2$ and λ is set to the average of the new value for λ_L and λ_U . In this case λ had been adjusted down as far as possible in 2d and a new feasible solution was not found, therefore the lower bound is increased.

It is possible that SP-A will converge at a local optimum that is inferior to the global optimum because of aggregation. Puelz (1999) discusses these conditions in detail relative to an

asset/liability management model with quadratic utility. Even in this case the SP-A model can be considered a portfolio improvement model that will derive a VaR-feasible, higher expected return portfolio given an initial portfolio plan.

References

- Bai, D., Carpenter, T., Mulvey, J., 1997. Making a case for robust optimization. *Management Science* 43, 895-907.
- Basak, S., Shapiro, A., 1999., Value-at-Risk Based Risk Management: Optimal Policies and Asset Prices. Working paper, The Wharton School.
- Beder, T.S. 1995., VAR: Seductive but dangerous. *Financial Analysts Journal* Sep/Oct, 12-24.
- Birge, J.R., Rosa, C.H., 1995. Modeling investment uncertainty in the costs of global CO2 emission policy. *European Journal of Operational Research* 83, 466-488.
- Cariño, D.R., Kent, T., Myers, D.H., Stacy, C., Sylvanus, M., Turner, A.L., Watanabe, K., Ziemba, W.T., 1994. The Russell-Yasuda Kasai model: An asset/liability model for Japanese insurance company using multistage stochastic programming. *Interfaces* 24, 29-49.
- Cariño, D.R., Ziemba, W.T., 1998. Formulation of the Russell-Yasuda Kasai financial planning model. *Operations Research* 46, 433-449.
- Dahl, H., Meeraus, A., Zenios, S.A., 1993. Some financial optimization models: I risk management. in *Financial Optimization*, Zenios (Ed.), Cambridge University Press, Cambridge, 3-36.
- Duffie, D., Pan, J., 1997. An overview of value at risk. *Journal of Derivatives* 4, 7-49.
- Golub, B., Holmer, M., McKendall, R., Pohlman, L., Zenios, S.A., 1995. A stochastic programming model for money management. *European Journal of Operational Research* 85, 282-296.

- Hiller, R.S., Eckstein, J., 1993. Stochastic dedication: Designing fixed income portfolios using massively parallel benders decomposition. *Management Science* 39, 1422-1438.
- Holmer, M.R., Zenios, S.A., 1995. The productivity of financial intermediation and the technology of financial product management. *Operations Research* 43, 970-982.
- Kalin, D., Zagst, R., 1999. Portfolio optimization: volatility constraints versus shortfall constraints. *OR Spektrum*. 21, 97-122.
- Koskosidis, Y.A., Duarte, A.M., 1997. A scenario-based approach to active asset allocation. *Journal of Portfolio Management* 23, 74-85.
- Lucas, A., Klaassen, P., 1998. Extreme returns, downside risk, and optimal asset allocation. *Journal of Portfolio Management* 25, 71-79.
- Markowitz, H.M. 1959., *Portfolio Selection: Efficient Diversification of Investments*, John Wiley, New York.
- McKay, R., Keefer, T.E., 1996. VaR is a dangerous technique. *Corporate Finance - Searching for systems integration supplement* 30.
- Mulvey, J.M., Vanderbei, R.J., Zenios, S.A., 1995. Robust optimization of large-scale systems. *Operations Research* 43, 264-281.
- Pritsker, M. 1997., Evaluating value at risk methodologies: Accuracy versus computational time. *Journal of Financial Services Research* 12, 201-242.
- Puelz, A.v. 1999., Stochastic convergence model for portfolio selection, Social Science Research Network (SSRN) working paper series (www.ssrn.com).
- Simons, K. 1996., Value at risk - new approaches to risk management. *New England Economic Review* Sep/Oct, 3-13.
- Stambaugh, F., 1996. Risk and value at risk. *European Management Journal*. 14, 612-621.

Uryasev, S., Rockafellar, R.T., 1999. Optimization of Conditional Value-at-Risk. Research Report #99-4, Center for Applied Optimization at the University of Florida.

Vladimirou, H., Zenios, S.A., 1997. Stochastic linear programs with restricted recourse. *European Journal of Operational Research* 101, 177-192.

Young, M.R., 1998. A minimax portfolio selection rule with linear programming solution. *Management Science* 44, 673-683.