Axioma

Flexible is better.®

Axioma Research Paper
No. 66

May 25, 2016

A CVaR Scenario-based Framework: Minimizing Downside Risk of Multi-asset Class Portfolios

Multi-asset class (MAC) portfolios can be comprised of investments in equities, fixed-income, commodities, foreign-exchange, credit, derivatives, and alternatives such as real-estate and private equity. The return for such *non-linear* portfolios is *asymmetric* with significant tail risk. The traditional Markowitz Mean-Variance Optimization (MVO) framework, that linearizes all the assets in the portfolio and uses the standard deviation of return as a measure of risk, does not accurately measure risk for such portfolios. We consider a scenario-based "Conditional Value-At-Risk" (CVaR) approach for minimizing the downside risk of an existing portfolio with MAC overlays. The approach uses (a) Monte Carlo simulations to generate the asset return scenarios, and (b) incorporates these return scenarios in a scenario-based convex optimization model to generate the overlay holdings. We illustrate the methodology on three examples in the paper: (1) hedging an equity portfolio with index puts; (2) hedging a callable bond portfolio with interest rate caps; (3) hedging the credit spread risk of a convertible bond portfolio. We compare the CVaR approach with parametric MVO approaches that linearize all the instruments in the MAC portfolio, and show that (a) CVaR approach generates portfolios with better downside risk statistics, (b) CVaR hedges return more attractive risk decompositions and stress test numbers—tools commonly used by risk managers to evaluate the quality of hedges.

A CVaR Scenario-based Framework: Minimizing Downside Risk of Multi-asset Class Portfolios

Kartik Sivaramakrishnan and Robert Stamicar

May 2016

1 Introduction

Multi-asset class (MAC) portfolios are an integral component of asset manager, asset owner and hedge fund investments and consist of a broad set of assets that include equities, fixed-income, commodities, foreign exchange, credit, and alternatives such as real estate and private equity. MAC instruments provide a more diversified set of asset allocation opportunities that are based on a wide spectrum of risk and return profiles. For example, adding investment grade bonds to an equity portfolio can improve its risk profile, while adding high yield corporates will provide greater returns at the cost of elevated risk; or securities tied to LIBOR can help dampen interest rate risk in a fixed-income portfolio.

With a broad set of assets, portfolio construction is challenging under a MAC setting. First, the characteristics of a portfolio should align with the investment process. Whether a top-down or bottom-up allocation process is utilized, or whether return over yield is chosen, these characteristics need to be translated to security selection. In addition, constraints including manager's preferences and institutional mandates need to be enforced in the model. For example, liquidity constraints can limit investment in alternatives or emerging markets. Moreover, portfolio construction is often a combination of subjective and quantitative analysis. For instance, risk decompositions and exposures at the factor and asset levels should intuitively align with a given strategy. In Figure 1, we schematically, on the left hand side, represent the investment process. This figure is by no means exhaustive but serves to illustrate the complexity of portfolio construction for MAC portfolios.

We do not consider portfolio construction in the *traditional* sense where we adjust all portfolio holdings so as to optimize an objective function in this paper. Instead, we consider the risk management problem of mitigating tail risk from a *hedging* perspective. There has been a lot of interest in downside risk protection especially in the wake of the 2008 financial crisis. Our starting point is a *base* portfolio that has been constructed, say, from an investment process that may even involve the use of sophisticated tools such as an optimization algorithm. The downside risk of the existing portfolio is then minimized with the addition of hedging overlays over a specified time horizon.

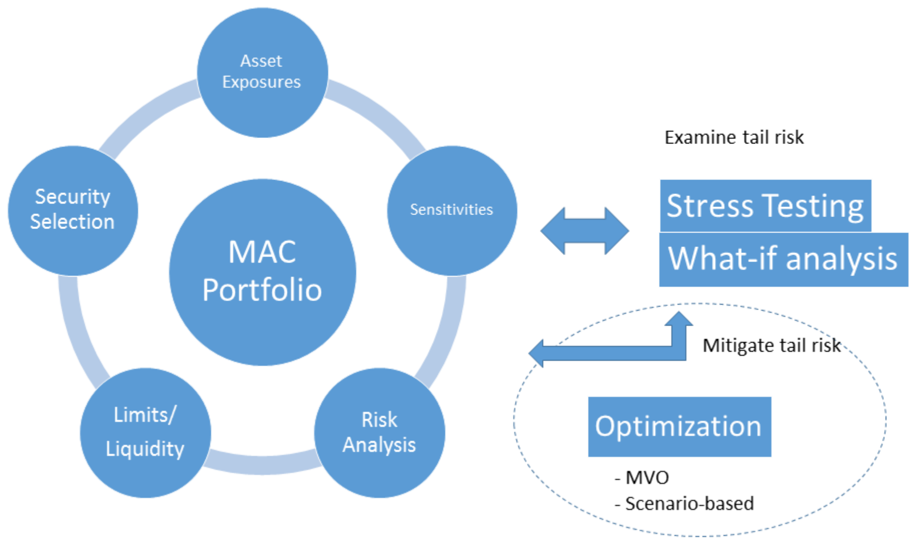

Figure 1: Moving parts in MAC portfolio construction

The Markowitz mean-variance optimization (MVO) (see Markowitz [10]) framework that is widely used in equity portfolio management is not suitable for MAC portfolios. In particular, positions with nonlinear payoffs introduce asymmetry between positive and negative returns and as a result can lead to significant skewness and kurtosis for the overall portfolio return distribution. The traditional MVO approach uses the standard deviation of returns as a risk measure. To illustrate how standard deviation understates the tail risk in a MAC portfolio, consider the upper and lower exhibits in Figure 2 that show the return distributions for a covered call portfolio (long 100 shares of S&P 500 and short one call contract on the S&P 500). The nonlinear payoff of the call introduces an asymmetry between positive and negative returns of the underlying S&P 500 index. A large negative return of the underlying may have little effect on a short call, but a large positive return will significantly decrease its value. The upside of the covered call portfolio is capped by the strike since the call will get exercised if the S&P 500 price exceeds the strike. The MVO model will not capture these nonlinear effects. In addition, parametric optimization approaches (see Jondeau & Rockinger [6]) that incorporate higher moments, such as skewness and kurtosis in the MVO framework, are also fairly limited in the size of MAC portfolios that they can handle. More accurate portfolio construction with MAC portfolios, therefore, requires that one look beyond the parametric MVO framework.

We present a scenario-based hedging framework to mitigate the tail risk for multi-asset class portfolios. The framework consists of two phases: In the first phase, we generate Monte Carlo simulations for asset and hedging position returns. In the second phase, the simulations from the first step are fed into a scenario-based convex optimization problem where tail risk is minimized by selecting overlay hedges, which are constrained by a budget. Although these techniques are well known in the literature (see for example Glasserman [3] for Monte Carlo portfolio simulations and Rockafellar & Uryasev [12, 13] for scenario-based CVaR optimization) it is important to highlight two key points. First, the simulation-based approach is only as good as the pricing and risk models that are used to generate asset returns. We will illustrate, later in one of our test cases, the pitfalls when an inappropriate pricing model is used under our framework. Second, it is important to acknowledge that the scenario-based approach, however useful, should not be a standalone analysis. It is a complement to existing risk management tools such as risk factor decomposition and what-if analysis. The base and hedged portfolios should be compared under these traditional risk management tools and we will provide examples throughout this paper.

The contribution of this paper is two-fold:

- We use a novel approach to regularize the CVaR optimization problem and a unique methodology to ensure that we have enough scenarios in this model. The model is solved with the interior point optimizer.
- We highlight the CVaR hedging approach on three diverse MAC portfolios. For each example, we compare the CVaR approach with a parametric MVO approach that linearizes all the instruments in the portfolio.

The paper is organized as follows: Section 2 introduces downside risk measures and describes VaR and CVaR in detail. Section 3 describes the two phases of the scenario-based CVaR approach. Sections 4, 5, 6 apply the CVaR hedging approach to minimize the downside risk of three diverse MAC portfolios. Section 7 presents our conclusions.

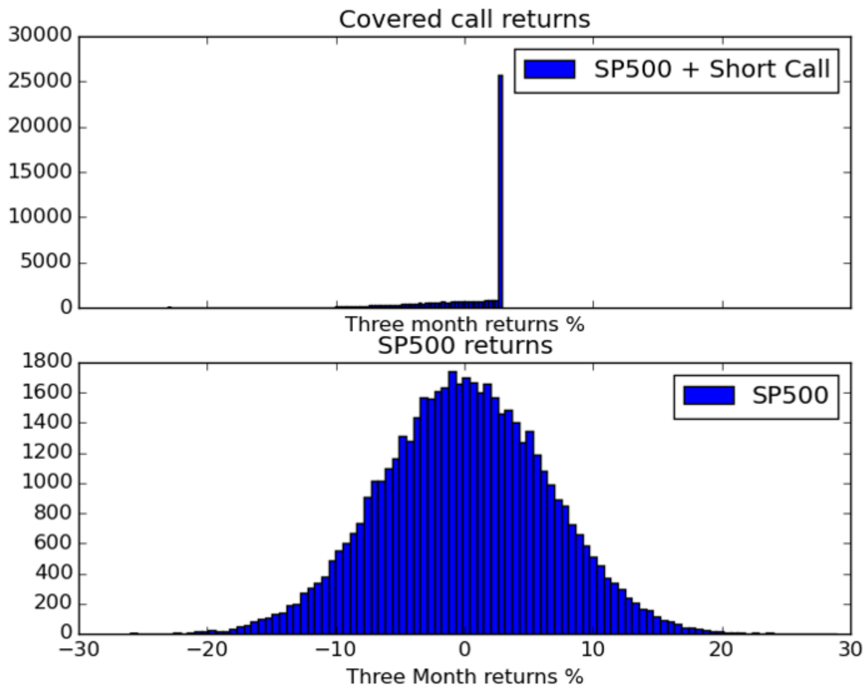

Figure 2: Covered Call portfolio

2 Downside risk measures

There has been a lot of interest in downside risk measures that reflect the financial risk associated with losses in a MAC portfolio. The most popular downside risk measures include Value at Risk (VaR) and Conditional Value at Risk (CVaR), also known as Expected Shortfall. VaR played a prominent role in the Basel regulatory framework (see McNeil et al. [11]) and is widely used as a risk measure for MAC portfolios. However, it should be noted that CVaR will replace VaR for all risk and capital calculations under the Basel Commmittee's FRTB (fundamental review of the the trading book) and is expected to be implemented by 2018. This change in regulatory framework is due to CVaR overcoming several shortcomings of VaR.

VaR at confidence level ϵ is the $(1-\epsilon)$ percentile of the portfolio return distribution. VaR has a simple interpretation: A portfolio VaR at the 95% confidence level over a 10 day period of \$10 million implies that we are 95% confident that the portfolio will not suffer losses greater than \$10 million over a 10 day period. CVaR at confidence level ϵ is the expected value of the loss exceeding VaR, and it was introduced to overcome some of the following shortcomings of VaR:

- VaR has poor mathematical properties. It is not coherent, i.e., not sub-additive in the framework of Artzner et al. [2]. If VaR is used to set risk limits, it can lead to concentrated portfolios. CVaR, on the other hand, is a coherent risk measure encouraging diversification.

- VaR does not measure the left tail of the portfolio return distribution. The worst case loss can be much larger than VaR. CVaR, on the other hand, is a tail statistic that measures the losses that occur in the left tail of the return distribution.
- VaR for a non-linear portfolio is difficult to optimize in practice as it requires the solution to a non-convex optimization problem. CVaR, on the other hand, can be optimized via a scenario-based linear program—this follows from the seminal work of Rockafellar and Uryasev (see [12, 13]).

3 CVaR scenario-based framework

We describe the scenario-based CVaR framework in this section. There are two main parts to this framework: (a) Section 3.1 presents the Monte Carlo framework that generates the return scenarios for the different MAC instruments in the portfolio, (b) Section 3.2 presents a scenario-based convex optimization problem that takes the Monte Carlo scenarios as inputs and then generates the overlay hedges. More precisely, we keep the holdings of the original or base portfolio fixed and adjust the weights of hedging/overlay positions to minimize CVaR. Section 3.3 presents our methodology to ensure that we have enough scenarios in the CVaR optimization problem.

3.1 Phase 1: Monte Carlo framework to generate asset scenarios

We first describe the Monte Carlo framework that is used to generate the instrument scenarios. The asset returns are driven by a set of market factors. We will assume that the returns of the n assets are driven by k market factors. Let $\hat{f}$ denote the vector of factor returns representing either the relative or the absolute changes in the factor level values f. For factors such as equity prices and volatility, we define the factor return to be the relative or log change in the factor values. For fixed-income factors such as interest rates and credit spreads, we define the factor return to be the absolute change in the factor value.

The Monte Carlo framework that we outline below is fairly general and is based on the simulation of copulas. A given multivariate distribution of risk factor returns can be decomposed into marginal distributions, which describes the individual risk factors, and the copulas, which describe the *dependence* structure. A copula, C, is a distribution function with standard uniform marginal distributions u_i:

$$C(u) = C(u_1, \ldots, u_k) \tag{1}$$

We can combine the copula with marginal distributions, F_i, of risk factor returns to generate a multivariate distribution:

$$\begin{aligned} C(F_1(x_1),\ldots,F_n(x_n)) &= Pr[U_1 \le F_1(x_1),\ldots,U_n \le F_n(x_n)] \\ &= Pr[F_1^{-1}(U_1) \le x_1,\ldots,F_n^{-1}(U_n) \le x_n] \\ &= Pr[X_1 \le x_1,\ldots,X_n \le x_n] \\ &= F(x_1,\ldots,x_n) \end{aligned} \tag{2}$$

The simulation of the copula is performed on the uniform random variables U_i in (2) (see for example, McNeil et al. [11]). After we simulate the uniform random variables, we obtain the joint distribution of risk factor returns, $\hat{f} = X = (X_1,\ldots,X_n)$, by linking the marginal distribution to the copula via $X_i = F_i^{-1}(U_i)$. More precisely, we use a parametric k-variate joint distribution C^θ, such as a Gaussian or Student multivariate distribution, to model the dependence structure. The simulation of the copula and asset values is generated as follows:

Simulate parametric copula:

- Simulate $(Z_1,\ldots,Z_k)$ from C^θ
- Simulate parametric copula: $U_i = C_i^\theta(Z_i)$, where C_i^θ are the marginal distributions corresponding to C^θ

Simulate risk factors and asset values:

- Extract risk factors returns via marginals: $\hat{f}_i = X_i = F_i^{-1}(U_i)$
- Compute simulated factor levels from returns: $\hat{f}_i \mapsto f_i$
- Apply pricing models to compute simulated values for asset p: $MV_p = MV_p(f_1,\ldots,f_k)$

As an illustration consider Figure 3 which shows the Monte Carlo framework at work. The returns of stocks A and B have a correlation of 0.6 and follow a normal and student t-distribution, respectively. We use a Gaussian copula to generate the joint distribution of stock A and B returns in the top two exhibits of Figure 3. The bottom left exhibit provides a scatter plot of stock returns by applying the inverse marginal mapping to the simulated uniform random variables. The histogram plot along the y-axis shows that the marginal distribution of the stock B return is a student-t distribution with longer tails. The bottom right exhibit of Figure 3 shows a scatter plot of the price of stock A and the price of a long put on stock B. The histogram plot on the y-axis shows the distribution of the put prices. Notice how this histogram is asymmetric with a long right tail indicating that the long put has limited downside but significant upside.

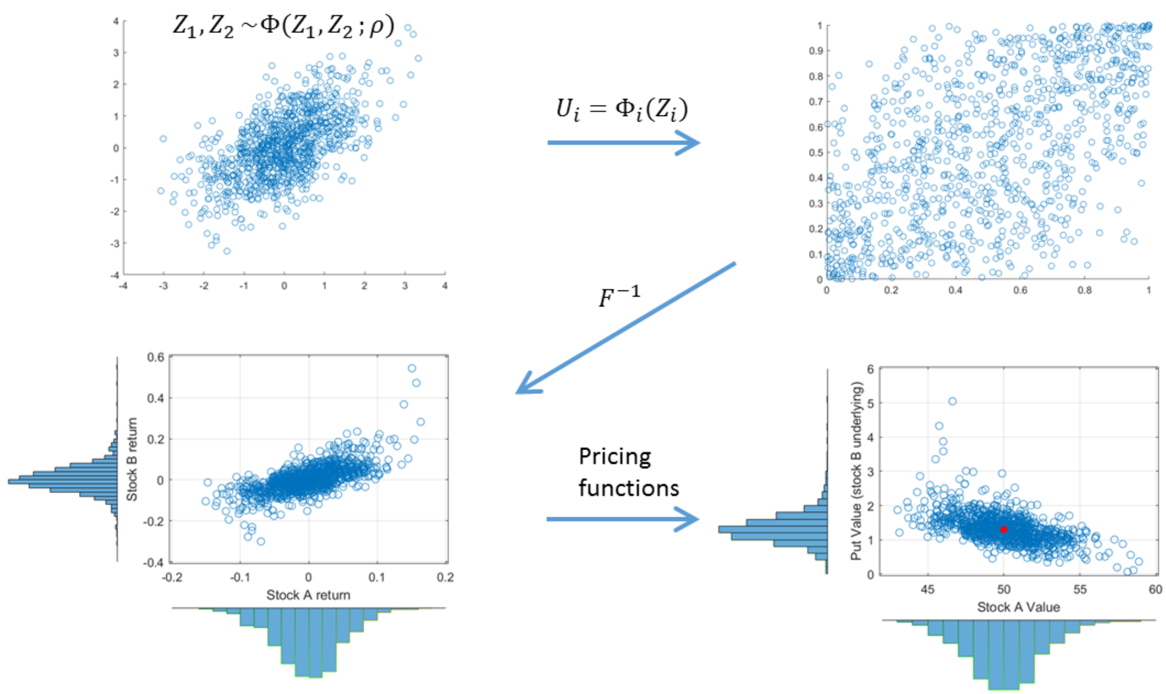

Figure 3: Monte Carlo framework at work

For completeness, we summarize the Monte Carlo algorithm for Gaussian and Student copulas below.

Algorithm 1 (Monte Carlo framework for generating asset scenarios)

1. *Simulate s joint factor return scenarios from Gaussian or a Student-t copula (see McNeil et al. [11]) with marginal normal and student-t distributions of the form*

$$C(u_1, \ldots, u_k) = \Phi^{\Sigma}\left(\Phi_1^{-1}(u_1), \ldots, \Phi_k^{-1}(u_K)\right)$$

where $\Phi^{\Sigma}()$ denotes the cumulative distribution function of a multivariate normal or a student-t distribution with covariance matrix Σ, each Φ_i^{-1}, $i = 1, \ldots, k$ denotes the inverse cumulative distribution function of an univariate standard normal or student-t distribution, and each u_i, $i = 1, \ldots, k$ denotes a uniform random number between 0 and 1.

2. *Compute the profit/loss (P/L) for asset p in scenario j as*

$$PL_p^j = MV_p(f^j) - MV_p(f_0) \tag{3}$$

where

   (a) MV_p is an appropriate pricing function for asset p,

   (b) f_0 is the current vector of factor values that drive the return of asset p, and

   (c)

$$f_i^j = f_{0,i} + \hat{f}_i^j \quad or \quad f_i^j = f_{0,i}(1 + \hat{f}_i^j)$$

   depending on whether the return $\hat{f}_i$ for factor f_i is defined in terms of a relative or absolute change. f^j denotes the vector of factor values in scenario j at the end of the hedging time horizon.

Let PL denote the $n \times s$ (asset by scenario) P/L matrix.

3. *Scale each row of PL by the holding size (number of positions) of the asset to generate an $n \times s$ matrix of asset returns R.*

A few comments are now in order for risk factors and pricing models:

Risk Factors: The mapping of risk to pricing factors is important since this will affect asset scenario generation. The framework presented above is flexible enough to incorporate various degrees of granularity. For example, consider an equity portfolio containing all the assets in the S&P 500. One can choose the 500 equity prices as a set of 500 (granular) risk factors. Alternatively, one can choose the factors in a linear factor model as risk factors and map these factors to the equity prices via the exposures in the model. For example, in the Axioma US fundamental model, there are only 78 fundamental factors that represent a *parsimonious* set of risk factors. We will use a parsimonious set of risk factors for all equity positions we consider in this paper.

Pricing Models: The choice of pricing models in Step 2 of Algorithm 1 is important and should reflect the inherent risk of a given strategy. We will provide an example in Section 6 that highlights that the scenario-based hedging optimization framework is only as good as the pricing models.

3.2 Phase 2: CVaR-based optimization

The asset scenarios are fed to a scenario-based convex optimization that generates the asset holdings in the second part of the CVaR framework. We briefly describe this formulation below.

Let $w = (w_1, w_2, \ldots, w_n)$ denote the market value dollar holdings in the different assets in the portfolio. Consider the following problem:

$$\min_{w \in \mathcal{C}} \quad \mathrm{CVaR}(w, \epsilon) \tag{4}$$

where ϵ is the desired confidence level and $\mathcal{C}$ contains all the constraints representing the managers' preferences and institutional mandates. Instead of minimizing the CVaR of the portfolio, one can impose a risk limit on the CVaR of the portfolio that is of the form

$$\mathrm{CVaR}(w, \epsilon) \quad \leq \quad \theta \tag{5}$$

where θ represents an appropriate upper bound. Rockafellar-Uryasev [12] show that the model (4) can be formulated as a scenario-based linear program (LP). The solution to this LP is sensitive to small variations in the scenarios (see Gotoh et al. [4] and Lim et al. [9]). In our implementation, we assume that the scenarios are not point estimates but rather that the ith scenario lies in the following ellipsoidal uncertainty set:

$$\mathcal{Z}^i \quad = \quad \{r^i : (r^i - \bar{r}^i)^T Q^{-1} (r^i - \bar{r}^i) \leq \kappa^2\}, \tag{6}$$

where Q is an appropriate risk model possibly in linear factor form. We use the covariance matrix obtained by linearizing all the instruments in the portfolio (see Section 4.2 for an illustration on an equity and options portfolio). To immunize the solution to the CVaR optimization problem, i.e., make it relatively insensitive to small changes in the scenarios, we solve the following robust CVaR problem:

$$\min_{w \in \mathcal{C}} \quad \max_{r^i \in \mathcal{Z}^i, i=1,\ldots,s} \quad \mathrm{CVaR}(w, \epsilon). \tag{7}$$

We show in the Technical Appendix at the end of the paper that the robust CVaR can be formulated as the following scenario-based, second-order cone program:

$$\begin{array}{lrcl} \min\limits_{w,\alpha,u} & \alpha + \frac{1}{s(1-\epsilon)} \sum\limits_{i=1}^{s} u_i + \kappa \sqrt{w^T Q w} & & \\ \text{s.t.} & (\bar{r}^i)^T w + \alpha + u_i & \geq & 0, \\ & u_i & \geq & 0, \quad i = 1, \ldots, s, \\ & w & \in & \mathcal{C} \end{array} \tag{8}$$

where $\bar{r}^i$ denotes the nominal scenario value in the i column of the asset-scenario matrix R, and $\mathcal{C}$ contains linear constraints on the asset holdings. This formulation is solved with an interior point optimizer. Note that:

- The only difference between (8) and the Rockafellar-Uryasev LP formulation is the presence of the risk term $\sqrt{w^T Q w}$ in the objective function.
- More scenarios are needed when the confidence level ϵ increases and techniques such as importance sampling (see Glasserman [3]) can be used to selectively generate more scenarios in the tail of the distribution.
- The number of auxiliary variables u is equal to the number of scenarios. These variables also appear in the constraints of the model. The auxiliary variable u_i measures the excess loss in the ith scenario over VaR. This variable is zero if the loss in the ith scenario is less than VaR implying that only scenarios with positive u_i actually contribute to the CVaR of the portfolio.
- The α variable provides an estimate for the VaR of the portfolio.
- Specialized first-order and decomposition approaches (see Kunzi-Bay & Meyer [8] and Iyengar & Ma [7]) are available to solve (8) approximately and quickly when the number of samples is large.

3.3 Sensitivity of CVaR optimization to number of samples

The CVaR optimization uses scenarios and one has to test the stability of the solution with respect to the number of scenarios. We want to ensure that we have enough scenarios in the optimization so the optimal CVaR values exhibit two different types of stability: (a) *in-sample* stability where the optimal CVaR values exhibit small variation across scenarios of the same size, (b) *out-of-sample* stability where the optimized portfolios constructed across scenarios of the same size exhibit small variation when their CVaR values are computed on a master set of scenarios. The master set contains a large number of scenarios and is assumed to represent the return distribution of the assets in the portfolio. We perform the following analysis:

1. Start with a master set of 50,000 scenarios.
2. Choose 1,000 subsamples of a given size (1,000; 5,000; and 10,000 say) from the master set.
3. For each of the 1,000 samples, compute the CVaR optimized portfolio.
4. Plot a histogram of the optimal objective values for these 1,000 optimizations when testing for *in-sample* stability.
5. Use the master set of 50,000 scenarios to compute a histogram of CVaR objective values for these optimized portfolios when testing for *out-of-sample* stability.

We can be confident that we have enough scenarios in the CVaR optimization if the variation of the CVaR values in the *in-sample* and *out-of-sample* tests is small. Figure 4 shows the variation in the optimal objective values with different number of samples in the CVaR optimization. There is little variation in the values when 10,000 samples are employed in the optimization. One can see that there is a more variation in the in-sample values—this indicates that we need more samples to accurately compute the CVaR value for an optimized portfolio than a static portfolio. We run a similar test for all the examples in this paper to ensure that we have enough scenarios in the CVaR optimization.

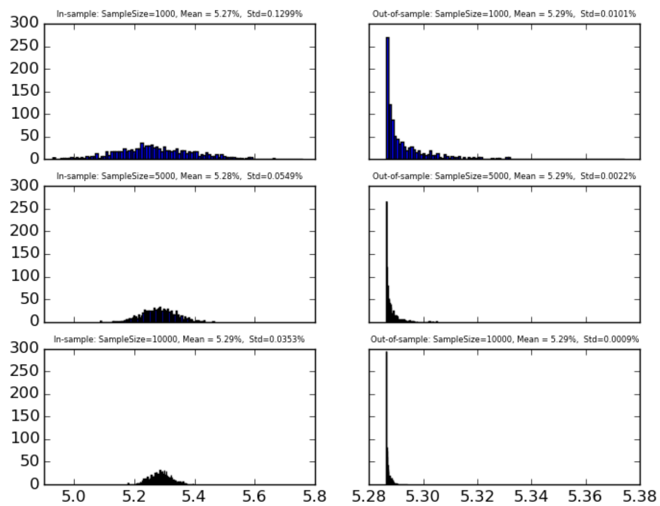

Figure 4: Variation in CVaR with number of samples in optimization

4 Hedging an equity portfolio with index options

We hedge an equity portfolio with index puts in this section. The equity portfolio contains all the assets in the S&P 500. Our strategy is given below:

1. Hedge a long-only equity portfolio on Dec 19, 2014 over a three-month horizon with SPX puts (European puts on the S&P 500) that have different strikes and three months to expiration.
2. Minimize the CVaR of the overall portfolio (equities + puts) at the 95% confidence level over the three-month hedging horizon.
3. Option budget constraint.
4. Cannot change the equity holdings.

5. Only purchase puts.

It is worth emphasizing that one can also give the CVaR approach the flexibility to write puts. This can be done to collect the premium from suitable out-of-money puts and presumably supplement the available budget, but we do not pursue put writing in this section. This section is organized as follows: We describe the CVaR approach for this example in Section 4.1. Two parametric MVO approaches that use a linear model for the options are described in Sections 4.2 and 4.3. We compare the CVaR approach with the MVO approaches in Section 4.4.

4.1 CVaR hedging

We first describe how we generate asset returns from the Monte Carlo simulation of risk factors.

1. Step 1 of Algorithm 1 first generates the joint distribution

$$\left(f^E, f^r, f^\sigma\right) \quad \in \quad N\left(0, \Sigma\right) \tag{9}$$

   of the fundamental equity factor, risk-free rate, and implied volatility returns that is assumed to follow a multivariate normal distribution $N(0, \Sigma)$ with mean zero and covariance Σ.

   (a) Risk-free rate (LIBOR) returns f^r include the one- and six-month factors. We use these factors to interpolate the three-month risk-free rate return that is used to price the puts.

   (b) Implied volatility returns f^σ from the SPX volatility surface include the six-month out-of-the-money, at-the-money, and in-the-money factors. We use the at-the-money factor to price the at-the-money puts, the out-of-the-money factor to price the out-of-the-money puts, and the in-the-money factor to price the in-the-money puts.

   The asset-specific risk returns are then simulated from a multivariate normal distribution with a diagonal covariance matrix Θ^2. The fundamental equity factor, risk-free rate, implied volatility, and the asset-specific returns represent our choice of risk factors for the Monte Carlo framework. Note that the asset-specific returns are uncorrelated with the other risk factors and also with each other.

2. Equity returns come from the Axioma equity fundamental factor model as

$$r^E \quad = \quad Bf^E + \epsilon \tag{10}$$

   where B contains the equity exposures to the fundamental factors, f^E contains the fundamental factor returns, and $\epsilon \in N\left(0, \Theta^2\right)$ contains the asset-specific returns where Θ^2 is a diagonal matrix containing the specific variances.

3. The SPX put returns are generated from the Black-Scholes pricing model:

$$r^O = \text{BS}(r^u, f^r, f^\sigma) \tag{11}$$

where $r^u = c^T r^E$ denotes the S&P 500 returns with c containing the S&P 500 weights on the rebalancing date. We illustrate the return computation for the ith SPX put in the jth scenario below. Let S and r denote the current level of the S&P 500 and the risk-free rate, respectively.

   (a) Let p_i, K_i, σ_i, and T_i denote the current price, strike, implied volatility, and the time to expiration of the ith option, respectively.

   (b) Let ΔT denote the length of the hedging time horizon.

   (c) Let $S^j = S(1 + r^{uj})$, $r^j = (r + f^{r,j})$, and $\sigma_i^j = \sigma_i(1 + f^{\sigma,j})$ denote the level of the S&P 500, risk-free rate, and the implied volatility in scenario j at the end of the rebalancing time horizon.

   (d) The Black-Scholes formula gives the price of ith SPX put in the jth scenario at the end of the rebalancing time horizon as

$$p_i^j = e^{-r^j(T_i - \Delta T)}\left(K_i \Phi(-d_2) - S^j \Phi(-d_1)\right) \tag{12}$$

   where

$$d_1 = \frac{\log\left(\frac{S^j}{K_i}\right) + \left(r^j + \frac{1}{2}\left(\sigma_i^j\right)^2\right)(T_i - \Delta T)}{\sigma_i^j \sqrt{T_i - \Delta T}}$$

$$d_2 = d_1 - \sigma_i^j \sqrt{T_i - \Delta T}$$

   and $\Phi()$ is the cumulative normal distribution function. The return of the ith put in the jth scenario is then given by

$$r_i^{O,j} = \left(\frac{p_i^j - p_i}{p_i}\right).$$

4.2 MVO delta-rho-vega hedging

Consider the following delta-rho-vega linear model (see Hull [5]):

$$r_i^O = \Delta_i \left(\frac{S_u}{p_i}\right) r^u + \left(\frac{\rho_i}{p_i}\right) f^r + \nu_i \left(\frac{\sigma_i}{p_i}\right) f^\sigma \tag{13}$$

for the ith option return, where p_i, Δ_i, ρ_i, σ_i, and ν_i represent the option share price, delta, rho, implied volatility, and vega, respectively; and S_u represents the current level of the S&P 500.

Assuming that the equity returns r^E are given by the linear factor model in equation (10), one can construct a linear model

$$r^p = Gf^p + H\epsilon \tag{14}$$

where $r^p = \begin{pmatrix} r^E \\ r^O \end{pmatrix}$ denotes the asset returns; $f^p = \begin{pmatrix} f^E \\ f^r \\ f^\sigma \end{pmatrix}$ denotes the set of fundamental, risk-free rate, and volatility surface factors; ϵ denotes the asset-specific risk factors; and G and H are matrices of appropriate dimensions. We assume that $f^p \in N(0, \Sigma)$, and $\epsilon \in N(0, \Delta^2)$. It follows from equation (14) that the asset returns follow a multivariate normal distribution whose covariance matrix is given by

$$Q = \left(G\Sigma G^T + H\Theta^2 H^T\right).$$

Given portfolio holdings $w^p = \begin{pmatrix} w^E \\ w^O \end{pmatrix}$, the risk of the portfolio is given by

$$\begin{aligned} \text{Portfolio risk} &= (w^p)^T Q w^p \\ &= (w^p)^T \left(G\Sigma G^T + H\Theta^2 H^T\right) w^p. \end{aligned} \tag{15}$$

The MVO delta-rho-vega approach minimizes the total risk of the portfolio subject to all the constraints in the hedging strategy. A few comments are in order:

- The MVO delta-rho-vega approach is linearizing the options in the portfolio. Over the short hedging horizon of three months, the option delta dominates the rho and the vega. In this setting, the MVO approach mimics the *delta hedging* approach commonly used by traders.
- The MVO delta-rho-vega approach disregards the higher option moments, such as the gamma, that are important for longer hedging horizons.
- Since we are hedging the long-only equity portfolio with several index puts, the MVO delta-rho-vega approach has *multiple* solutions since there are several ways to choose the holdings in the SPX puts to effectively make the portfolio delta-neutral. In Section 4.3, we show how one can incorporate some of the higher options moments in the hedge by choosing an optimal portfolio that minimizes the portfolio gamma.

4.3 Enhanced MVO delta-gamma hedging

The MVO delta-gamma approach tries to incorporate option gamma in the MVO delta-rho-vega based hedging approach. Given that the MVO delta-rho-vega model has multiple solutions, the MVO delta-gamma hedging chooses an MVO delta-rho-vega solution that minimizes the portfolio gamma.

The portfolio gamma is given by

$$\text{Portfolio gamma} \quad = \quad \sum_{i \in \mathcal{O}} \left(\frac{\Gamma_i}{p_i} \right) w_i \tag{16}$$

where $\mathcal{O}$ is the set of options in the portfolio, and Γ_i, p_i, and w_i denote the gamma, share price, and the dollar holding in the ith option, respectively. Note that only the options in the portfolio contribute to the portfolio gamma.

The MVO delta-gamma follows a two level hierarchical process to generate the option holdings:

1. In level 1, we follow the MVO delta-rho-vega approach, where we minimize the portfolio risk given by equation (15) subject to all the constraints in the hedging strategy.
2. In level 2, we minimize the portfolio gamma given by equation (16) subject to all the constraints in Level 1 and an additional constraint prescribing that

$$\begin{aligned} \text{Portfolio risk} \quad &= \quad (w^p)^T \left(G \Sigma G^T + H \Theta^2 H^T \right) w^p \\ &\leq \quad \text{optimal objective value from Level 1 plus tolerance.} \end{aligned}$$

The solution to the Level 2 problem gives the option hedges.

4.4 Comparing the CVaR and MVO hedges

We compare the CVaR, MVO delta-rho-vega, and MVO delta-gamma optimized portfolios for different option budgets. The upper and lower exhibits in Figure 5 plot the histogram of 50,000 realizations of the CVaR, delta-rho-vega, delta-gamma, and the unhedged portfolio returns (in that order) at the end of the three-month rebalancing horizon, when the option budget is 1% and 5% of the reference size of the equity portfolio, respectively. The downside risk statistics for the different portfolios for a 5% budget are summarized in Table 1. When the budget is small, the three hedges give similar results. However, when the budget is large, Table 1 shows that the CVaR portfolio return has the best downside risk statistics: The worst loss, CVaR, VaR, and standard deviation are much smaller than those of the MVO delta-rho-vega and MVO delta-gamma portfolios, respectively. Moreover, the right CVaR for the CVaR portfolio is only slightly smaller than that of the two MVO portfolios and the unhedged portfolio showing that the CVaR portfolio has the best skew among these portfolios, i.e., better downside without taking much away from the upside. Although the MVO delta-rho-vega objective values for the MVO delta-rho-vega and MVO delta-gamma portfolios are smaller than the corresponding value for the CVaR hedged portfolio (not surprisingly, since these portfolios minimize this objective function), the downside risk statistics are worse. The poor performance of the MVO approaches can be attributed to the linear model for the puts that does not capture the asymmetric payoffs of these instruments.

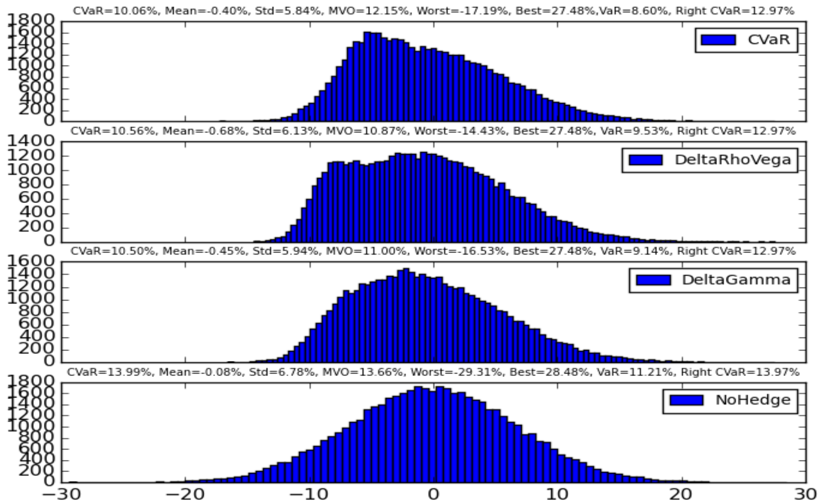

(a) Budget 1%

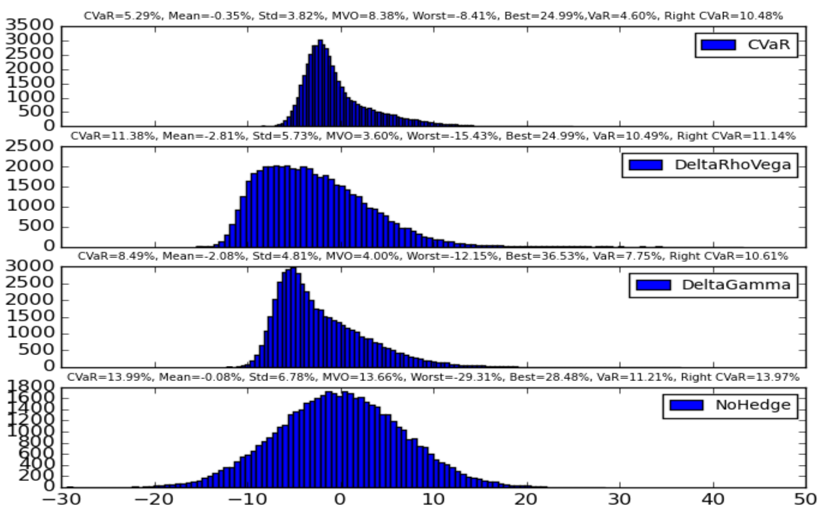

(b) Budget 5%

Figure 5: Example 1: Histogram of future portfolio returns

Statistic	CVaR Scenario-based	DeltaRhoVega	DeltaGamma	Unhedged
CVaR	5.29%	11.38%	8.49%	13.99%
VaR	4.60%	10.49%	7.75%	11.21%
Worst Loss	-8.41%	-15.43%	-12.15%	-29.31%
StdDev	3.82%	5.73%	4.81%	6.78%
Right CVaR	10.48%	11.14%	10.61%	13.97%
MVO Objective	8.38%	3.60%	4.00%	13.66%

Table 1: Example 1: Downside Risk Statistics Budget 5%

The upper and lower exhibits of Figure 6 compare the strikes of the puts chosen by the hedging approaches for option budgets of 1% and 5%, respectively. To give a perspective, the S&P 500 was trading at 2070.65 on Dec 19, 2014. When budget is 5%, the CVaR approach purchases the expensive put with strike 2075 that offers the best downside protection. The MVO delta-rho-vega approach purchases the less expensive put with strike 1925 that offers less downside protection. Neither approach uses its entire budget. The MVO delta-gamma is the only approach that uses up its entire budget—it spends about 60% of the budget in purchasing the less expensive put with strike 1925 and the rest of the budget in purchasing the most expensive put with strike 2075.

5 Example 2: Callable bond portfolio with IR caps

We consider a fixed-income portfolio with 48 US callable bonds on May 2, 2006. Figure 7 shows that most of the risk in this portfolio is interest rate risk. We want to hedge this portfolio against an increase in interest rates with OTC interest rate caps in this section. Callable bonds have fat left tails since their upside is capped when interest rates decrease. This is because the issuer of the callable bond may call, i.e., redeem the bond when its value increases beyond its call price. Our strategy is as follows:

1. Hedge callable bond portfolio with OTC interest rate caps.
2. There are 10 OTC interest rate caps with varying strikes and times to expiration:
   (a) Floating rate is based on the six-month LIBOR rate. The prevailing six-month LIBOR rate on May 2, 2006 is 5.34%.
   (b) Strikes are chosen to the prevailing six-month LIBOR rate.
   (c) Short-term caps expire on June 30, 2008 and long-term caps expire on June 30, 2010.
   (d) Each caplet tenor is six months.

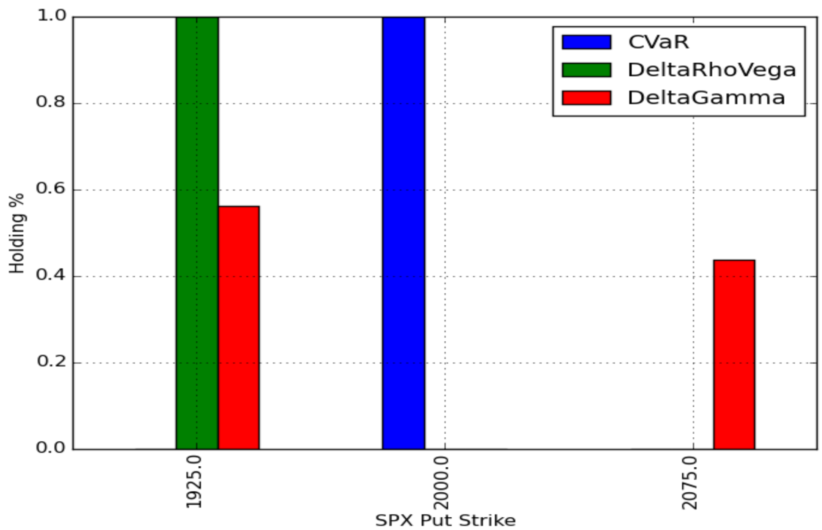

(a) Budget 1%

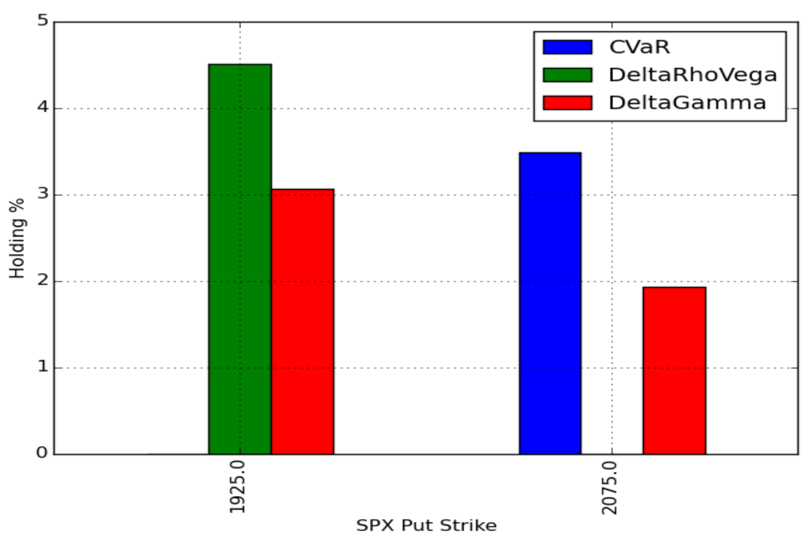

(b) Budget 5%

Figure 6: Example 1: Overlays purchased by portfolios

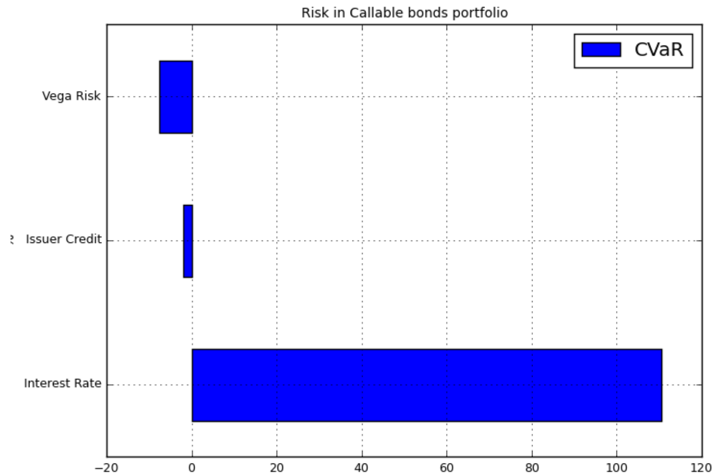

Figure 7: Sources of risk in callable bonds portfolio

3. Cannot change the callable bond holdings.
4. Experiment with budgets of 1% and 5% of the total market value of the callable bond portfolio.
5. Compare the CVaR and MVO hedges over a six-month hedging horizon.

We briefly describe the Monte Carlo pricer used to generate the asset scenarios below:

1. The Hull-White pricer (see Hull [5]) is used to price the callable bonds where the pricing factors include the US sovereign, swap, issuer credit, and swaption volatility factors.
2. Black's formula (see Hull [5]) is used to price the caps as a sequence of caplets where the pricing factors include the US sovereign, swap, issuer credit, and the cap volatility surface factors.

The MVO approach linearizes both the callable bonds (the underlying instruments) and the IR caps (hedging overlays) to arrive at a covariance matrix for the portfolio.

The upper and lower exhibits in Figure 8 plot the histogram of 20,000 realizations of the CVaR and the MVO hedged portfolio returns at the end of the six-month hedging horizon, when the option budget is 1% and 5% of the reference size of the callable bond portfolio, respectively. The histogram of the unhedged callable bond portfolio returns for these realizations is shown at the bottom of these two exhibits for comparison. Notice that the unhedged callable bond portfolio has a fat left tail since the upside of this portfolio is capped when the bonds are redeemed by the issuer. Table 2 presents the downside risk statistics for the different portfolios. Notice that the CVaR portfolio has superior downside risk statistics

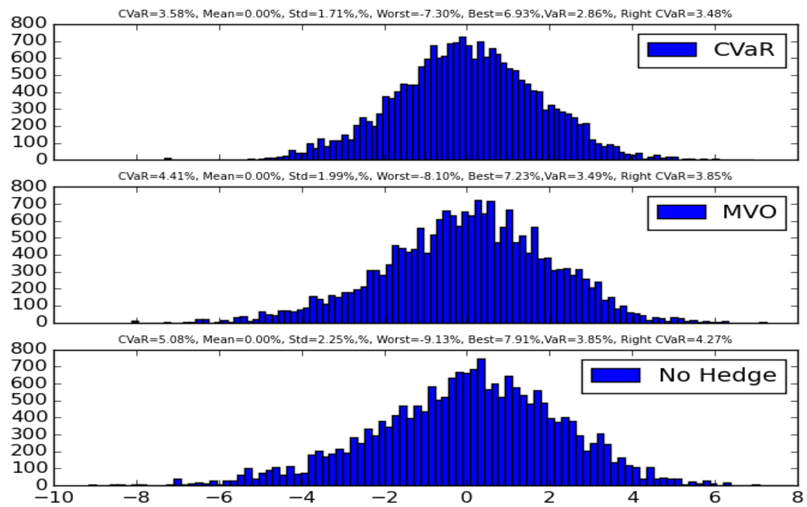

(a) Budget 1%

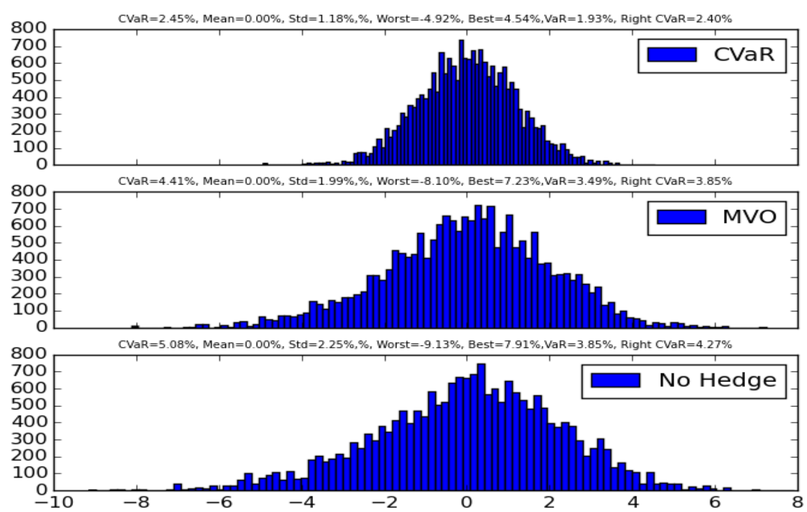

(b) Budget 5%

Figure 8: Example 2: Histogram of future portfolio returns

even for the small 1% budget: The worst-case return, left CVaR, left VaR, and standard deviation are much smaller than that of the MVO portfolio. The MVO approach also does not use up its entire 1% budget and so the results are unchanged when one increases the budget to 5%. On the other hand, the downside risk statistics for the CVaR portfolio only get better with the increased 5% budget.

Statistic	CVaR 1% Scenario-based	CVaR 5% Scenario-based	MVO	Unhedged
CVaR	3.58%	2.45%	4.41%	5.08%
VaR	2.86%	1.93%	3.49%	3.85%
Worst Loss	-7.30%	-4.92%	-8.10%	-9.13%
StdDev	1.71%	1.18%	1.99%	2.25%
Right CVaR	3.48%	2.40%	3.85%	4.27%

Table 2: Example 2: Downside risk statistics

The upper exhibit of Figure 9 plots the incremental CVaR decompositions for the different components of the CVaR portfolios with 1% and 2% budgets and the MVO portfolio. The incremental contribution for the overlays is the difference between the portfolio CVaR and the CVaR of the base portfolio, i.e., the portfolio with the overlays removed. The incremental contributions for the overlays are negative for the three portfolios indicating that the overlays are reducing the CVaR of the overall portfolio. Note that the greatest reduction in risk is for the CVaR hedge with 2% budget. We used 10,000 scenarios to compute these contributions. The lower exhibit in Figure 9 plots the marginal CVaR decompositions for the different components of the three portfolios. The marginal CVaR of the overlay is the change in the portfolio CVaR resulting from marginal changes in the overlay positions. Note that the marginal contribution for the overlay for the CVaR hedge with 2% budget is the smallest indicating that the portfolio CVaR is relatively insensitive, i.e., changes only slightly for small changes in the overlay positions. This is not surprising since this portfolio has the smallest CVaR among the three portfolios.

The upper exhibit of Figure 10 plots the P/Ls for the CVaR, MVO, and unhedged portfolios when the portfolios are subject to instantaneous interest rate shocks ranging from -1% to 1%. An interest rate shock of 1% implies that we move the entire yield curve up by 1%. When a shock of 1% is applied, the base portfolio loses 6%, the MVO portfolio loses 4%, the CVaR portfolio with 1% budget loses 2%, and the CVaR portfolio with 5% budget loses nothing. On the other hand, when the entire yield curve moves down by 1%, the base portfolio gains 5%, the MVO and the CVaR portfolios with 1% budget gain 4%, and the CVaR portfolio with 5% budget gains 2%. Notice that the CVaR portfolio with 1% budget has a smaller downside than the MVO portfolio while these portfolios have the same upside.

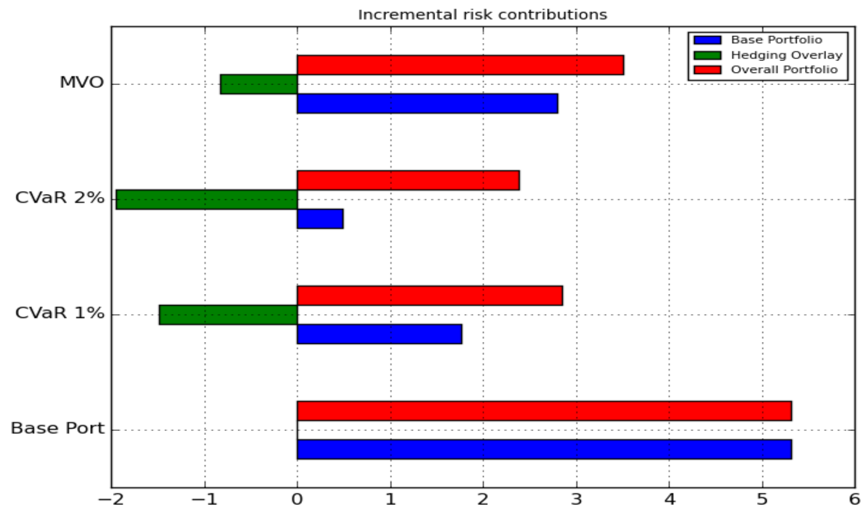

(a) Incremental Risk Contributions

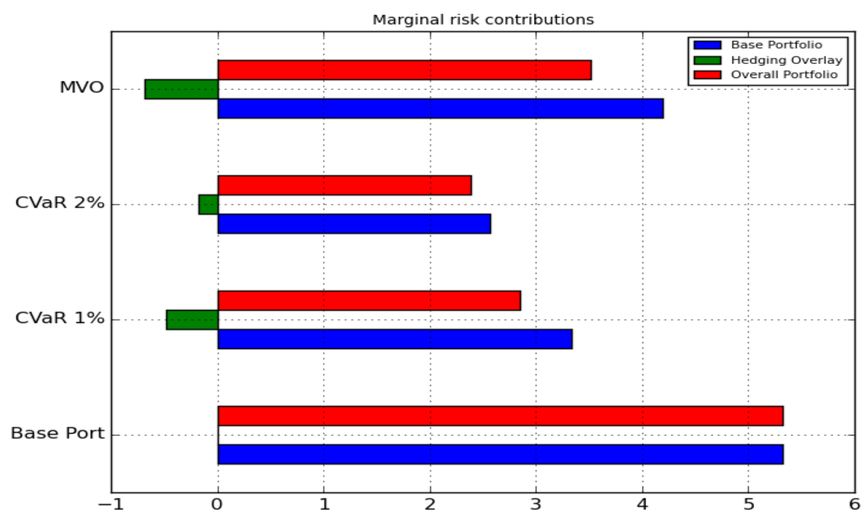

(b) Marginal Risk Contributions

Figure 9: CVaR Risk Decomposition for portfolios

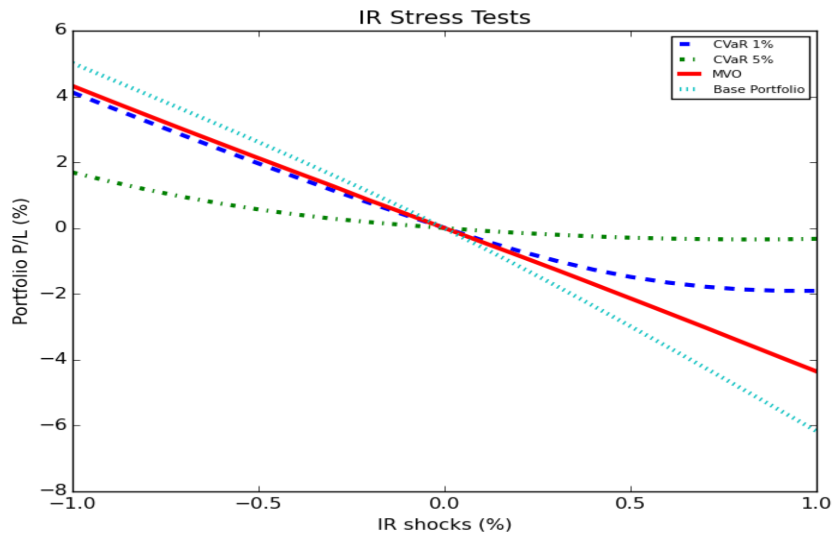

(a) Interest Rate Stress Tests

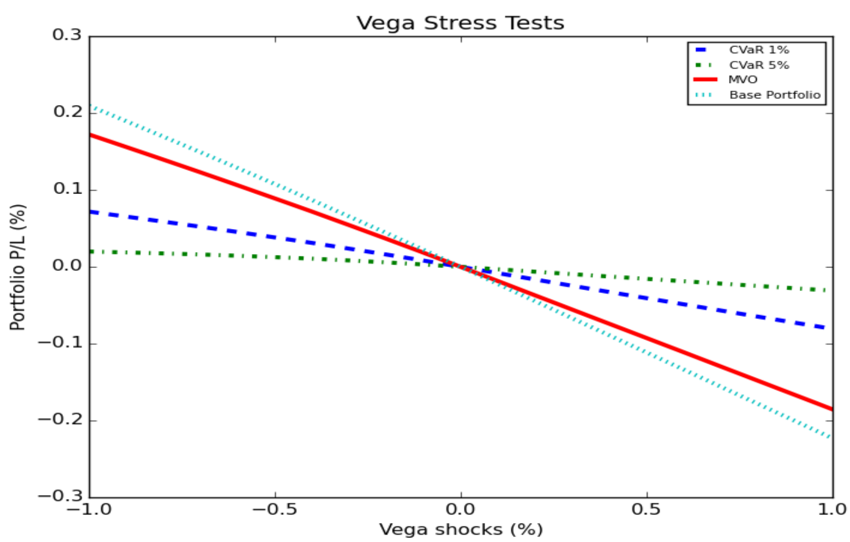

(b) Vega Stress Tests

Figure 10: Instantaneous Stress Tests

6 Example 3: Convertible bond portfolio

In this example, we highlight the importance of utilizing appropriate pricing models that capture the key risks for a given strategy. We consider a portfolio which is comprised of stocks and convertible bonds. For this portfolio, we compare two different pricing models for convertible bonds where one model captures credit spread risk, whereas the other does not.

Before specifying the details of the portfolio (and overlays), we briefly comment about the modeling of convertible bonds. A convertible bond is a hybrid instrument that contains both bond and equity-like features. An investor has the option to convert the bond into a number of predefined shares, specified by the conversion ratio. The conversion feature gives convertible bonds an embedded call option that increases when the underlying equity grows. For example, consider a convertible bond with a notional of €1000 and a conversion ratio of 25. If the underlying equity price rises to €50, then the likelihood of conversion is high since its intrinsic call value is significantly greater than its notional ($25 \times 50 = 1{,}250 > 1{,}000$).

In addition to the embedded call option, converts can contain other sources of optionality such as call/put schedules, contingent conversion thresholds and soft call provisions.[1] Figure 11 presents two different equity profiles for a convertible bond. The dotted line represents the parity price of the convertible; for large stock prices, we expect the investor to convert into underlying shares. The red line represents the *typical* payoff, which shows the asymmetric returns from the embedded call option, namely, a large positive stock return will increase the value of a convert significantly, whereas a large negative stock return will have little impact on its return. Here, the investor still enjoys the bond component even though the upside from the equity component has been dampened. However, this scenario of floored protection from negative stock returns is unrealistic. As the stock value decreases significantly, we expect the creditworthiness of the issuer to deteriorate. The solid blue line in Figure 11 depicts this case.

From a modeling perspective, there are various ways of collapsing the bond floor as the issuing firm's credit deteriorates. For example, Andersen and Buffum [1] introduce a jump-diffusion model and list common parameterizations of the intensity process that are monotonically decreasing functions of the stock price. Thus, as the stock price drops, the default intensity increases, which in turn results in the collapse of the bond floor. In addition to default intensity or hazard rate models, structural (risk-neutral) models will also provide a link from equity to credit. As the equity value deteriorates, the (risk-neutral) probability of default increase since the value of the firm's assets have decreased. In either modeling choice listed above, a more realistic equity profile is the solid blue line in Figure 11. For this section, we will model the credit risk of the convertible bonds utilizing a risk-neutral

[1] Soft calls allow the issuer to redeem the bond at par only when the stock price is above a certain threshold (this feature benefits the investor). Under contingent conversion conditions, investors can convert to underlying shares only when the stock price is above a specified threshold (this feature benefits the issuer).

structural model.[2]

Now we turn to the composition of the base portfolio and the hedging overlays:

- The base portfolio under consideration is comprised of stocks and convertible bonds. The stocks are constituents of the Euro Stoxx 50, Stoxx Oil & Gas, PSI. The portfolio is long convertible bonds with short positions in the underlying stock; the convertible bonds are in the energy sector.
- The overlay consists of an iTraxx Euro credit default index and a single name put that references a convert for a Portuguese issuer.

The credit index is used to mitigate risk in the equity portfolio. Our intuition is that the credit index is negatively correlated to the long equity positions, and, as a result, should have a negative risk contribution when added to the base portfolio. However, there are also short equity positions in the portfolio (from delta-hedging the convertible bonds), so it is unclear how much of the credit index might be selected. On the other hand, it is fairly clear that the single name put should be selected. Intuitively, as we minimize CVaR, we expect that the optimizer will select the put option since losses in the convertible bond will be offset by gains in the put position.

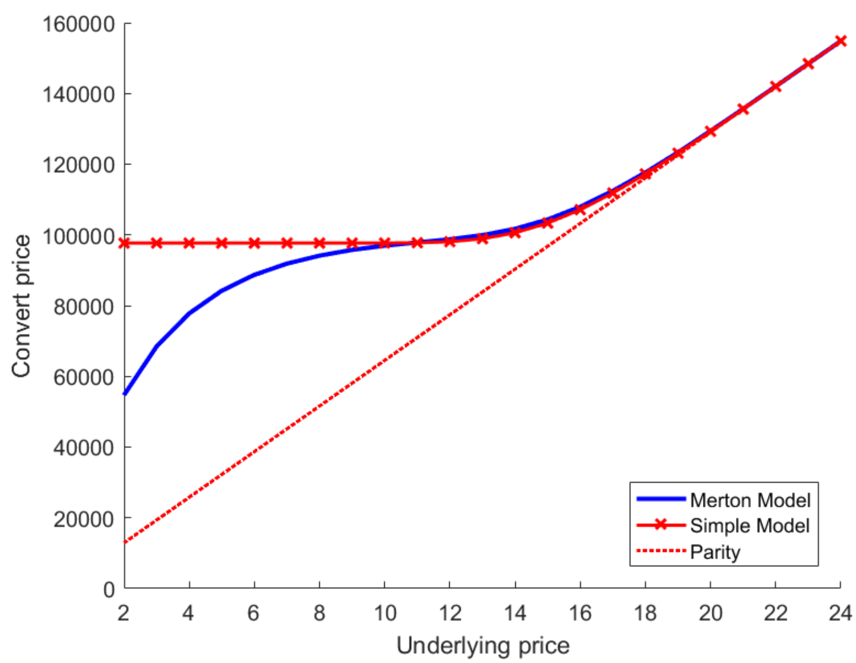

Figure 11: Convertible bond payoff

We will compare the overlays chosen by the CVaR approach when we use: (a) Merton model

[2] In fact, we combine the structural model with a traditional equity lattice pricer.

(blue line in Figure 11), and (b) simple pricing model (red line in Figure 11) that simply prices the convertible bond as a bond with an embedded equity option without any credit spread risk. The left exhibit of Figure 12 shows the histogram of the different portfolios at the end of the one-month hedging horizon. The right exhibit of Figure 12 shows the overlays purchased by the CVaR portfolios with the simple and Merton convertible pricers for a hedging budget of 1%. Note that the CVaR portfolio with the simple pricing model does not purchase any overlays while the CVaR portfolio with the Merton pricing model uses about 70% of its budget to purchase put protection. This is not surprising since only the Merton pricing model captures the credit spread risk of the convertible bond. The CVaR optimization engine in Phase 2 of the hedging approach is only as good as the input scenarios and care should be taken to choose appropriate pricing models that capture the length of hedging horizons, instruments characteristics, and the prevailing market conditions.

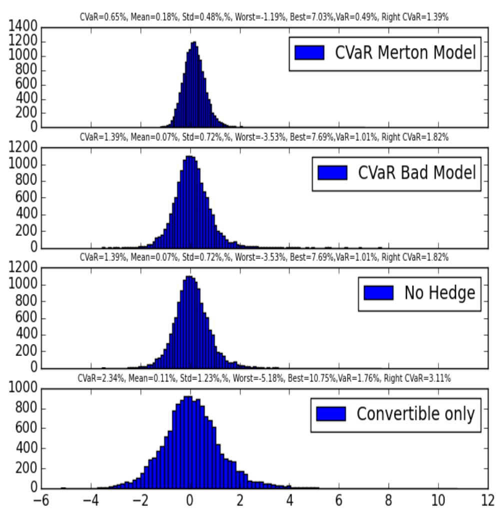

(a) Example 3: Histogram of future portfolio returns

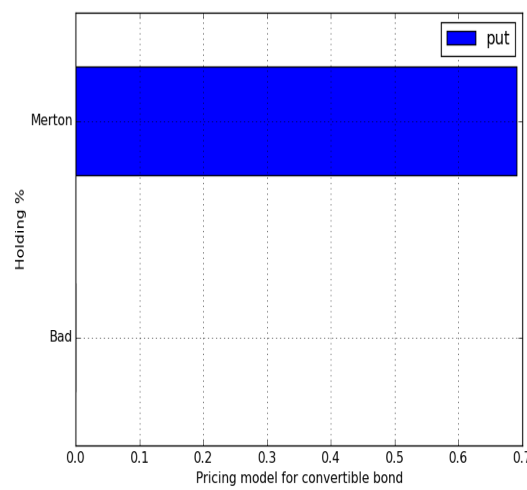

(b) Example 3: Overlays purchased

Figure 12: Example 3: Compare portfolios

7 Conclusions

We presented a two-phase scenario-based CVaR hedging approach for minimizing the downside risk in MAC portfolios. The base portfolio is fixed and the hedging approach determines the overlay holdings by minimizing the CVaR of the overall portfolio. The first phase of the hedging approach uses a Monte Carlo framework for generating the scenarios for the different instruments in the portfolio. The second phase incorporates these scenarios in a scenario-based convex optimization problem to generate the overlay holdings. We compare the CVaR approach with MVO-based approaches that linearize all the instruments in the portfolio on three examples and show that (a) the CVaR portfolio has better downside risk statistics, and (b) the CVaR portfolio returns more attractive risk decompositions and stress test numbers—tools that are commonly used by risk managers to evaluate the quality of hedges.

We must emphasize that the CVaR hedging approach is flexible; this includes (a) choice of the risk factors and the pricing models in the Monte Carlo framework, and (b) the setup of the scenario-based convex optimization, including how it is regularized to return stable optimal solutions. Ultimately, the CVaR hedging approach is not a standalone analysis and the flexibility should be used wisely and in conjunction with other tools (risk decomposition, delta-hedging, stress testing) in the arsenal of the risk manager.

References

[1] L. Andersen and D.B. Buffum (2003). Calibration and Implementation of Convertible Bond Models, *Journal of Computational Finance*, 7, 1–34.

[2] P. Artzner, F. Delbaen, J.M. Eber, and D. Heath (1999). Coherent Measures of Risk, *Mathematical Finance*, 9(3), 203–228.

[3] P. Glasserman (2003). *Monte Carlo Methods in Financial Engineering*, Springer.

[4] J. Gotoh, K. Shinozaki, and A. Takeda (2013). Robust Portfolio Techniques for Mitigating the Fragility of CVaR Minimization and Generalization to Coherent Risk Measures, *Quantitative Finance*, 13, 1621–1635.

[5] J.C. Hull (2008). *Options, Futures, and Other Derivatives*, 7th edition, Prentice Hall.

[6] E. Jondeau and M. Rockinger (2006). Optimal Portfolio Allocation under Higher Moments, *European Financial Management*, 12(1), 29–55.

[7] G. Iyengar and A.K.C. Ma (2013). Fast Gradient Descent Rule for Mean-CVaR Optimization, *Annals of Operations Research*, 205, 203–212.

[8] A. Kunzi-Bay and J. Mayer (2006). Computational Aspects of Minimizing Conditional Value-at-Risk, *Computational Management Science*, 3, 3–27.

[9] A. Lim, J.G. Shanthikumar, and G.Y. Yan (2011). Conditional Value-at-Risk in Portfolio Optimization: Coherent but Fragile, *Operations Research Letters*, 39, 163–171.

[10] H. Markowitz (1959). *Portfolio Selection: Efficient Diversification of Instruments*, Wiley.

[11] A.J. McNeil, R. Frey, and P. Embrechts (2015). *Quantitative Risk Management: Concepts, Techniques, and Tools*, Revised Edition, Princeton University Press.

[12] R.T. Rockafellar and S. Uryasev (2000). Optimization of Conditional Value-at-Risk, *Journal of Risk*, 2, 493–517.

[13] R.T. Rockafellar and S. Uryasev (2002). Conditional Value-at-Risk for General Loss Distributions, *Journal of Banking & Finance*, 26, 1443–1471.

A Robust CVaR formulation

Consider the robust CVaR problem

$$\min_{w \in \mathcal{C}} \max_{r^i \in \mathcal{Z}^i, i=1,\ldots,s} \quad \text{CVaR}(w, \epsilon) \tag{17}$$

where

$$\mathcal{Z}^i \quad = \quad \{r^i : ||Q^{-1/2}(r^i - \bar{r}^i)|| \leq \kappa\}. \tag{18}$$

Let $\bar{\alpha}$ be the VaR estimate variable. Define the auxiliary scenario variables

$$u_i \quad \geq \quad \max_{r^i \in \mathcal{Z}^i} \max\{(-r^i)^T w - \bar{\alpha}, \;\; 0\}, \;\; i = 1, \ldots, s. \tag{19}$$

The constraints that define these scenario variables can also be written as

$$\begin{aligned} \min_{r^i \in \mathcal{Z}^i} \; (r^i)^T w + \bar{\alpha} + u_i \quad &\geq \quad 0, \\ u_i \quad &\geq \quad 0, \;\; i = 1, \ldots, s. \end{aligned}$$

Using the alternative definition of the auxiliary scenario variables, we rewrite (17) as

$$\begin{array}{lrcl} \min\limits_{w, \bar{\alpha}, u} & \bar{\alpha} + \dfrac{1}{s(1-\epsilon)} \displaystyle\sum_{i=1}^{s} u_i & & \\ \text{s.t.} & \min\limits_{r^i \in \mathcal{Z}^i} \; (r^i)^T w + \bar{\alpha} + u_i & \geq & 0, \\ & u_i & \geq & 0, \;\; i = 1, \ldots, s, \\ & w & \in & \mathcal{C}. \end{array} \tag{20}$$

Minimizing a linear function over the ellipsoidal uncertainty set (18) can be done in closed form. We have

$$\min_{r^i \in \mathcal{Z}^i} (r^i)^T w \quad = \quad (\bar{r}^i)^T w - \kappa \sqrt{w^T Q w}, \;\; i = 1, \ldots, s. \tag{21}$$

Plugging in the scenario equations (21) into the model (20) and introducing a new variable

$$\alpha \quad = \quad \left(\bar{\alpha} - \kappa \sqrt{w^T Q w}\right)$$

helps us write the model (20) as

$$\begin{array}{lrcl} \min\limits_{w, \alpha, u} & \alpha + \dfrac{1}{s(1-\epsilon)} \displaystyle\sum_{i=1}^{s} u_i + \kappa \sqrt{w^T Q w} & & \\ \text{s.t.} & (\bar{r}^i)^T w + \alpha + u_i & \geq & 0, \\ & u_i & \geq & 0, \;\; i = 1, \ldots, s, \\ & w & \in & \mathcal{C}. \end{array} \tag{22}$$

This shows that the robust CVaR problem has an additional risk term in the objective function and is a scenario-based, second-order cone program.

Axioma

Flexible is better.®

United States and Canada: 212-991-4500

Europe: +44 20 7856 2451

Asia: +852-8203-2790

New York Office
Axioma, Inc.
17 State Street
Suite 2700
New York, NY 10004

Phone: 212-991-4500
Fax: 212-991-4539

Atlanta Office
Axioma, Inc.
400 Northridge Road
Suite 850
Atlanta, GA 30350

Phone: 678-672-5400
Fax: 678-672-5401

San Francisco Office
Axioma, Inc.
201 Mission Street
Suite #2150
San Francisco, CA 94105

Phone: 415-614-4170
Fax: 415-614-4169

London Office
Axioma, (UK) Ltd.
30 Crown Place
London, EC2A 4E8

Phone: +44 (0) 20 7856 2451
Fax: +44 (0) 20 3006 8747

Hong Kong Office
Axioma, (HK) Ltd.
Unit B, 17/F
30 Queen's Road Central
Hong Kong

Phone: +852-8203-2790
Fax: +852-8203-2774

Singapore Office
Axioma, (Asia) Pte Ltd.
30 Raffles Place
#23-00 Chevron House
Singapore 048622

Phone: +65 6233 6835
Fax: +65 6233 6891

Geneva Office
Axioma CH
Rue du Rhone 69, 2nd Floor
1207 Geneva, Switzerland

Phone: +41 22 700 83 00

Sales: sales@axioma.com
Client Support: support@axioma.com
Careers: careers@axioma.com