4

Mean-Variance portfolio allocation with a value at risk constraint[1]

Enrique Sentana
(CEMFI, CEPR and LSE FMG)
CEMFI. Casado del Alisal, 5. 28014 MADRID.
E-mail: sentana@cemfi.es

Abstract: In this paper, I first provide a simple unifying approach to static Mean-Variance analysis and Value at Risk, which highlights their similarities and differences. Then I use it to explain how fund managers can take investment decisions within the well-known Mean-Variance allocation framework that satisfy the VaR restrictions imposed on them by regulators. I do so by introducing a new type of line to the usual mean – standard deviation diagram, called *IsoVaR*, which represents all the portfolios that share the same VaR for a fixed probability level. Finally, I analyse the "shadow cost" of a VaR constraint.

1. INTRODUCTION

Mean-Variance analysis and Value at Risk (VaR) are tools routinely used by fund managers all over the world in their portfolio allocation decision-making. Therefore, it is not surprising that vast amounts of academic papers and books have been written on both topics, some of which discuss their relationship (see e.g. Alexander and Baptista (2000), Cuoco et al. (2001), Dowd (1999), Gourieroux and Monfort (2001), or section 14.4 of Jorion (1997)). Nevertheless, there is not yet a simultaneous treatment of the two that is "backward compatible" with the most widely used textbooks on financial investments. In this respect, the purpose of this paper is twofold. First, to offer a simple unifying approach, which highlights the similarities and differences between static Mean-Variance analysis and Value at Risk. And second, to explain how fund managers can take informed investment decisions that satisfy the risk-taking restrictions

often imposed on them by internal or external regulators without abandoning the well-known mean-variance allocation framework.

I do so by adding a new type of line to the usual mean - standard deviation diagram, which represents all the portfolios that share the same VaR (as a proportion of the available funds) for a fixed probability level. For obvious reasons, I call such lines *IsoVaRs*. Given some pre-specified risk control parameters, the corresponding IsoVaR effectively constrains the investment opportunity set, which is otherwise characterised by the usual portfolio frontier. Then, I simply have to superimpose the investors' risk preferences to find out the optimal restricted portfolio.

In this respect, this paper is more closely related to Gourieroux and Monfort (2001), who study efficient portfolios for parametric expected utility functions with general risk level constraints, and Cuoco et al. (2001), who analyse dynamic portfolio choice with a VaR limit when asset returns follow a multivariate geometric Brownian motion, than to Leibowitz and Kogelman (1991) and Alexander and Baptista (2000), who effectively replace mean-variance preferences by mean-VaR preferences. The concept of IsoVaR, though, turns out to be equivalent to the "shortfall line" introduced by Leibowitz and Kogelman (1991), although their emphasis was on achieving a minimum positive target return over the long run rather than on controlling losses over the short run.[2]

For simplicity of exposition, I make the assumption that a riskless asset is available for unlimited borrowing and lending at a common rate, although one could cope with the alternative assumptions of different rates, no borrowing, or no riskless asset with little additional effort. I also assume throughout that asset returns follow a joint multivariate elliptically symmetric distribution, of which the multivariate normal and the multivariate t are rather special cases. Apart from the obvious fact that many empirical studies indicate that the marginal distribution of asset returns has substantially fatter tails than the normal, and their cross-sectional dependence structure is non-linear, the two main advantages of weakening the standard Gaussianity assumption in this direction are that (i) mean-variance analysis still coincides with expected utility optimisation regardless of the exact form of the preferences, and (ii) the IsoVaRs remain straight lines in mean-standard deviation space.

The rest of the paper is organised as follows. I introduce the three basic building blocks (i.e. portfolio frontiers, indifference curves and IsoVaR lines) in section 2. Then, in section 3 I solve the general mean-variance portfolio allocation with a VaR cap problem, and analyse the "shadow cost" of the constraint. Finally, section 4 contains a brief summary of the results.

---

[2] For a recent, comprehensive analysis of multiperiod investment decisions, see Campbell and Viceira (2001).

Mean-Variance portfolio allocation...

2. THE THREE BUILDING BLOCKS

2.1 Mean-Variance Portfolio Frontiers

Consider a world with one riskless asset, and a finite number N of risky assets. Let R_S denote the gross return on the safe asset (that is, the total payoff per unit invested), $\mathbf{R} = (R_1, R_2, ..., R_N)'$ the vector of gross returns on the N remaining assets, with vector of means and matrix of variances and covariances $\boldsymbol{\upsilon}$ and $\boldsymbol{\Sigma}$ respectively, which I assume bounded. Let $p = w_S R_S + ... + w_N R_N$ denote the payoffs to a portfolio of the $N+1$ primitive assets with weights given by w_S and the vector $\mathbf{w} = (w_1, w_2, ..., w_N)'$. Importantly, I assume that there are no transaction costs or other impediments to trade, and in particular, that short-sales are allowed. I also assume that the wealth of any particular investor is such that her individual behaviour does not alter the distribution of returns.

There are at least three characteristics of portfolios in which investors are usually interested: their cost, the expected value of their payoffs, and their variance, given by $C(p) = w_S + w'\boldsymbol{\iota}$, $E(p) = w_S R_S + \mathbf{w'}\boldsymbol{\upsilon}$ and $V(p) = \mathbf{w'}\boldsymbol{\Sigma}\mathbf{w}$ respectively, where $\boldsymbol{\iota}$ is a vector of N ones. Let P be the set of payoffs from all possible portfolios of the $N+1$ original assets, i.e. the linear span of $(R_S, \mathbf{R'})'$. Within this set, several subsets deserve special attention. For instance, it is worth considering all unit cost portfolios R = $\{p \in P : C(p) = 1\}$ whose payoffs can be directly understood as returns per unit invested; and also all zero cost, or arbitrage portfolios A = $\{p \in P : C(p) = 0\}$. In this sense, note that any non-arbitrage portfolio can be transformed into a unit-cost portfolio by simply scaling its weights by its cost. Similarly, if $\mathbf{r} = \mathbf{R}-R_S\boldsymbol{\iota}$ denotes the vector of returns on the N primitive risky assets in excess of the riskless asset, it is clear that A coincides with the linear span of $\mathbf{r}$. The main advantage of working with excess returns is that their expected values $\boldsymbol{\mu} = \boldsymbol{\upsilon}-R_S\boldsymbol{\iota}$ directly give us the risk premia of $\mathbf{R}$, without altering their covariance structure. On the other hand, one must distinguish between riskless portfolios, S = $\{p \in P : V(p) = 0\}$ and the rest. In what follows, I shall impose restrictions on the elements of S so that there are no riskless "arbitrage"' opportunities. In particular, I shall assume that $\boldsymbol{\Sigma}$ is regular, so that S is limited to the linear span of R_S, and the law of one price holds (i.e. portfolios with the same payoffs have the same cost). I shall also assume that R_S is strictly positive (in practice, $R_S \geq 1$ for nominal returns).

A simple, yet generally incomplete method of describing the choice set of an agent is in terms of the mean and variance of all the portfolios that she can afford. Let us consider initially the case of an agent who has no wealth whatsoever, which means that she can only choose portfolios in A. In this context, frontier arbitrage portfolios, in the usual mean-variance sense, will be those that solve the program min $V(p)$ subject to the restrictions $C(p) = 0$ and $E(p) = \bar{\mu}$, with $\bar{\mu}$ real. Given that $C(p) = 0$ is equivalent to $p = \mathbf{w'r}$, I can re-write this problem as $\min_{\mathbf{w}} \mathbf{w'}\boldsymbol{\Sigma}\mathbf{w}$ subject to $\mathbf{w'}\boldsymbol{\mu} = \bar{\mu}$. There are two possibilities: (i) $\boldsymbol{\mu} = \mathbf{0}$, when the frontier can only be defined for $\bar{\mu} = 0$; or (ii) $\boldsymbol{\mu} \neq \mathbf{0}$, in which case the solution for each $\bar{\mu}$ is

$$\mathrm{w}^*(\bar{\mu}) = \bar{\mu}\,(\boldsymbol{\mu}\,\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu})^{-1}\,\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$$

As a consequence, the arbitrage portfolio $r_p = (\boldsymbol{\mu} \, \boldsymbol{\Sigma^{-1}} \boldsymbol{\mu})^{-1} \, \boldsymbol{\Sigma^{-1}} \mathbf{r}$ generates the whole zero-cost frontier, in what can be called one-fund spanning. Moreover, given that the variance of the frontier portfolios with mean $\bar{\mu}$ will be $\bar{\mu}^{\,2}(\boldsymbol{\mu}'\boldsymbol{\Sigma^{-1}}\boldsymbol{\mu})^{-1}$, in mean-standard deviation space the frontier is a straight line reflected in the origin whose efficient section has slope $\sqrt{\mu'\Sigma^{-1}\mu}$. Therefore, this slope fully characterises in mean-variance terms the investment opportunity set of an investor with no wealth, as it implicitly measures the trade-off between risk and return that the available assets allow at the aggregate level.

Traditionally, however, the frontier is usually obtained for unit-cost portfolios, and not for arbitrage portfolios. Nevertheless, given that the payoffs of any portfolio in R can be replicated by means of a unit of the safe asset and a portfolio in A, in mean-standard deviation space, the frontier for R is simply the frontier for A shifted upwards in parallel by the amount R_S. And although now we will have two-fund spanning, for a given safe rate, the slope $\sqrt{\mu'\Sigma^{-1}\mu}$ continues to fully characterise the investment opportunity set of an agent with positive wealth.

An alternative graphical interpretation of the same result would be as follows. The trade-off between risk and return of any unit-cost portfolio in R is usually measured as the ratio of its risk premium to its standard deviation. More formally, if $R_u \in$ R, where then $s_u = \mu_u / \sigma_u$, where $\mu_u = E(r_u)$, $\sigma_u = V(r_u)$, and $r_u = R_u - R_S$. This expression, known as the Sharpe ratio of the portfolio, remains constant for any portfolio whose mean excess return and standard deviation lie along the ray which, starting at the origin, passes through the point (μ_u, σ_u). As can be seen in Figure 1, geometrically the Sharpe ratio coincides with the slope of the appropriate ray. As a result, the steeper (flatter) a ray is (i.e. the closer to the *y (x)* axis), the higher (lower) the corresponding Sharpe ratio.

Figure 1
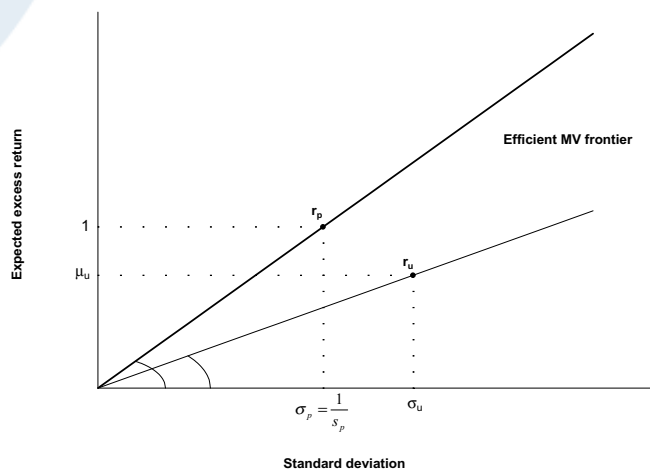
Then, since $\mu_p = 1$ and $\sigma_u^2 = (\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu})^{-1}$, the slope $s_p = \mu_u/\sigma_u = \sqrt{\boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}$ will give us the Sharpe ratio of

$$R_p(w_{r_p}) = R_s + w_{r_p}$$

for any $w_{r_p} > 0$, which is the highest attainable. Therefore, in mean excess return-standard deviation space, all $R_p(w_{r_p})$ lie on a positively sloped straight line that starts from the origin, as depicted in Figure 1. As the investor moves away from the origin, where she is holding all her wealth in the safe asset, the net total position invested in the riskless asset is steadily decreasing, and eventually becomes zero. Beyond that point, she begins to borrow in the money market to lever up her position in the financial markets.

2.2 Mean-Variance Indifference Curves

The analysis in the previous section shows that any investor endowed with unit wealth who only cares about the mean and variance of her returns will always choose some portfolio $R_p(w_{r_p})$, which exclusively combines R_s and r_p. Given my assumption that the joint distribution of the primitive asset returns belongs to the elliptical class,[3] such mean-variance preferences are compatible with standard Von-Neumann-Morgensten expected utility theory regardless of the specific functional form describing the investor's risk attitudes towards the underlying random payoffs (see Chamberlain (1983) or chapter 4, appendix B in Ingersoll (1987)). As a consequence, her optimal investment strategy will display two-fund separation. Nevertheless, the proposed allocation is incomplete, as the investor still has to determine the degree of leverage of her chosen position, w_{r_p}.

The traditional solution to such a scale indeterminacy is to use a particular parametrisation for the derived utility function that depends exclusively on the mean and variance of the portfolio returns. For instance, a commonly used specification is:

$$E(R_u) - \frac{\gamma}{2} V(R_u) = R_s + E(r_u) - \frac{\gamma}{2} V(r_u) \qquad (1)$$

---

[3] The vector of risky asset returns **R** will be distributed as an elliptically symmetric multivariate random variable of dimension N if and only if it can be written as an affine transformation of $\rho_t \cdot \mathbf{u}_t$, where $\mathbf{u}_t$ is uniformly distributed on the unit sphere surface in R^N, and ρ_t is a non-negative random variable that is independent of $\mathbf{u}_t$. When ρ_t is proportional to the square root of an F random variable with N and $1/\eta$ degrees of freedom, we obtain the multivariate t distribution, which in turn converges to the multivariate normal as $\eta \rightarrow 0$ (see Fang et al., 1990).

where γ is a positive risk aversion parameter.[4] The indifference curves associated with such preferences will be parallel straight half-lines with slope $\gamma/2$ in mean-variance space, and parallel half parabolas in mean-standard deviation space, as depicted in Figure 2. The indifference curves become steeper as the risk aversion coefficient γ increases. In particular, they will be almost vertical for an extremely risk averse individual, while they would be flat for a risk neutral agent (i.e. $\gamma = 0$). For given risk preferences, though, the utility level of the agent increases as we move up in a northwesterly direction across her indifference curves.
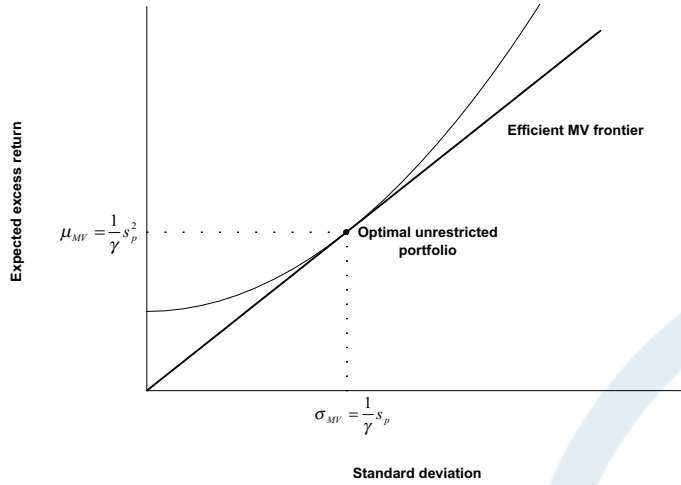
Figure 2
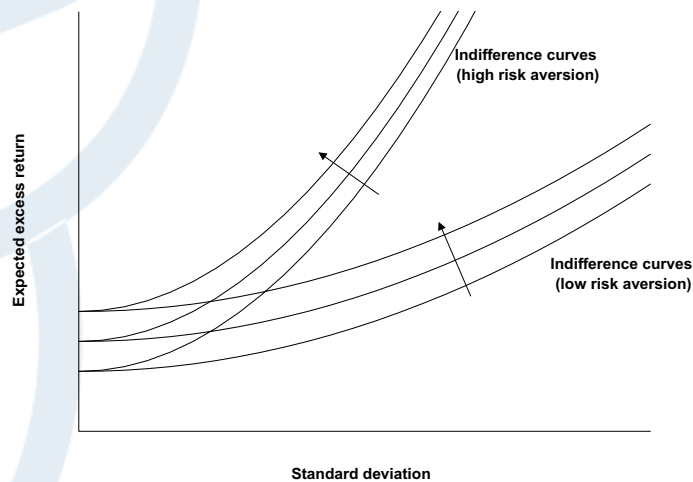
Therefore, it is not surprising that in the absence of any further constraints, her preferred portfolio will be such that the slope of the indifference curve at that point coincides with the maximum Sharpe ratio attainable, as can be seen in Figure 3. For the preferences in (1), in particular, this will happen for $r_u = \mathrm{w}_{r_p}^{\mathrm{MV}} r_p$, where

$$\mathrm{w}_{r_p}^{\mathrm{MV}} = \frac{1}{\gamma} \frac{E(r_p)}{V(r_p)} = \frac{1}{\gamma} (\mu' \Sigma^{-1} \mu) = \frac{1}{\gamma} \mathrm{s}_p^2$$

---

[4] The derived utility function in (1) is usually justified by a combination of constant absolute risk aversion and multivariate normality, which are unnecesarily restrictive assumptions for the purposes of this paper. Nevertheless, it helps illustrate the most important features of more general mean-variance preferences.

Mean-Variance portfolio allocation...

Figure 3

Figure showing curves labeled "Efficient MV frontier", "Optimal unrestricted portfolio", with axes "Expected excess return" (vertical) and "Standard deviation" (horizontal), and annotations $\mu_{MV} = \frac{1}{\gamma} s_p^2$ and $\sigma_{MV} = \frac{1}{\gamma} s_p$.

2.3 IsoVaRs

Nowadays, however, many institutional investors regularly use risk management procedures based on the ubiquitous Value at Risk to control for the risks associated with their portfolios.

Let $W > 0$ denote the initial wealth of an investor. Hence, if $R_u \in R$ denotes the gross return on her chosen portfolio, the random final value of her wealth will be

$$WR_u = WR_S + Wr_u$$

which contains both a safe component, WR_S, and a random component Wr_u. Therefore, the probability that she suffers a reduction in wealth larger than some fixed positive threshold value V will be given by the following expression

$$P[W(1 - R_S) - Wr_u \geq V] = P(r_u \leq 1 - R_S - V/W)$$

$$= P\left[\left(\frac{r_u - \mu_u}{\sigma_u}\right) \leq \frac{1 - R_s - V/W - \mu_u}{\sigma_u}\right] = F\left[\frac{1 - R_s - V/W - \mu_u}{\sigma_u}\right]$$

where μ_u and σ_u^2 are the mean and variance of r_u, and $F(.)$ is the cumulative distribution function of a zero mean – unit variance random variable within the appropriate elliptical class.[5] The value of V which makes the above probability equal to some pre-specified value α ($0 < \alpha < 1/2$)

---

[5] Due to the properties of the elliptical distributions (see theorem 2.16 in Fang et al, 1990), the cumulative distribution function $F(.)$ does not depend in any way on μ, Σ or the vector of portfolio weights w_u. In contrast, $F(.)$ will usually depend on some "tail thickness" parameters (e.g. the degrees of freedom in the multivariate t). In any case, note that the tail behaviour of the distribution of portfolio returns can be significantly non-Gaussian.

is known as the $100\alpha\%$ Value at Risk of the portfolio R_u. For convenience, though, the portfolio VaR is often reported in fractional form as *V/W*.

As we saw before (cf. Figure 1), there are infinitely many portfolios in R that will share the same Sharpe ratio, and they all line in a ray starting from the origin in mean excess return - standard deviation space. It turns out there are also infinitely many portfolios in R that will share the same fractional VaR *V/W* for a fixed probability level α. In fact, it is straightforward to see all such portfolios will also lie along a straight line in mean excess return-standard deviation space. Specifically,

$$\mu_u = (1 - R_S - V/W) - F^{-1}(\alpha) \cdot \sigma_u \tag{2}$$

where $F^{-1}(.)$ is the quantile (or inverse cumulative distribution) function of a zero mean - unit variance variate within the corresponding elliptical class. For obvious reasons, we shall call these locuses *IsoVaR* lines, or simply *IsoVaRs*.

Since we have assumed that $0 < \alpha < 1/2$, the IsoVaRs will be positively sloped, the more so the smaller α is. In fact, they will become almost vertical as $\alpha \to 0$, and almost flat as $\alpha \to 1/2$, as depicted in Figure 4. In contrast, changes in the target fractional VaR, *V/W*, will simply produce parallel shifts in the IsoVaRs. In particular, increases (decreases) in *V/W* will produce southeastward (northwestward) movements in the *IsoVaRs*, reflecting the fact that the VaR restriction has become looser (tighter).

Figure 4
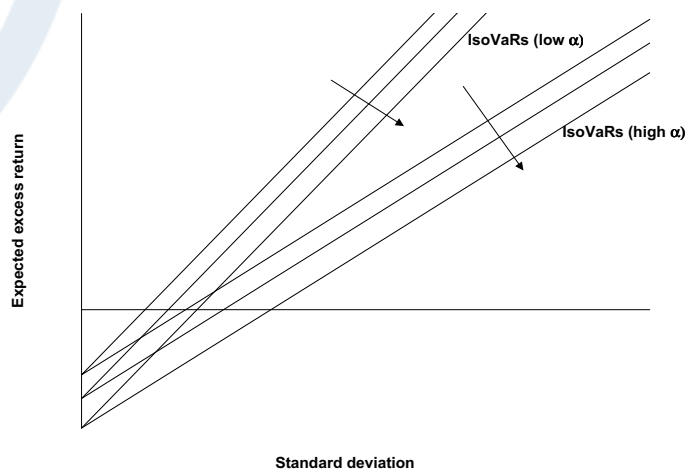
3. MEAN-VARIANCE PORTFOLIO ALLOCATION WITH A VAR CAP

As we saw in section 2.1, an investor endowed with unit wealth who only cares about the mean and variance of her returns will always choose some portfolio R_p (w_{r_p}), which exclu-

Mean-Variance portfolio allocation...

sively combines R_s and r_p. Suppose, however, that because of an internal or external regulation, the investor's overriding concern is that she can only risk losing a fraction *V/W* or larger of her initial wealth with probability α. Then, her effective choice set will be given by the section of the efficient mean-variance frontier between the origin and its intersection with the IsoVaR corresponding to α and *V/W*. Therefore, her optimal investment strategy will still display two-fund separation, the only effect of the VaR constraint being a reallocation of funds from r_p to R_s.[6]

Given that along the efficient frontier $E[r_p(w_{r_p})] = w_{r_p}$ and $V[r_p(w_{r_p})] = w_{r_p} \sigma_p$, I can solve for w_{r_p} from (2) to find that the intersection point will be:

$$w^I_{r_p} = \frac{(1 - R_s - V/W)}{1 + F^{-1}(\alpha)/s_p}$$

as long as the VaR significance level α is such that $1 + F^{-1}(\alpha)/s_p < 0$ for $R_s - 1 + V/W \geq 0$. In this respect, note that if α is high enough so that $-F^{-1}(\alpha) \leq s_p$, then $w^I_{r_p}$ becomes unbounded. The reason is that the overall risk-return trade-off that the fund manager faces is so good that the portfolio VaR restriction is effectively void.

Therefore, the optimal constrained portfolio for an agent with mean-variance preferences will be the minimum of $w^{MV}_{r_p}$ and $w^I_{r_p}$. If the agent is sufficiently risk averse, then $w^{MV}_{r_p}$ will also solve the restricted problem, while if the agent is less risk averse, then $w^I_{r_p}$ will be chosen instead. The optimal solution in the latter case is represented in Figure 5.

Figure 5
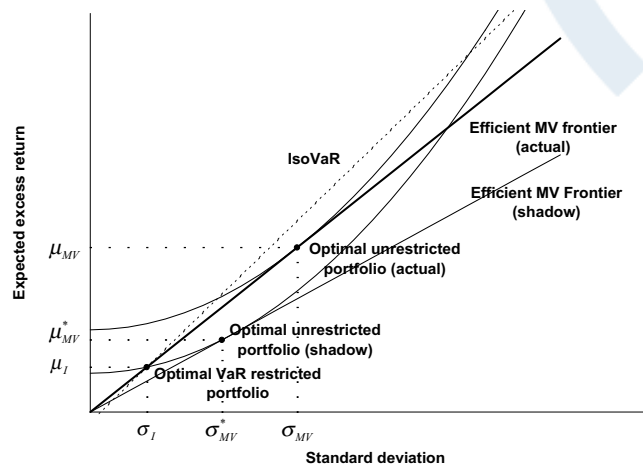
---
[6] The same result is obtained by Cuoco et al. (2001) in their dynamic set-up.

The previous analysis highlights an important fact that is very often blurred in discussions of the effects of VaR constraints on portfolio allocation. As this figure clearly illustrates, the expected return of the constrained portfolio cannot be larger than the expected return of the unconstrained portfolio, because the fund manager must put aside some additional funds in the safe asset to achieve the required VaR when the restriction is binding. Not surprisingly, the risk of the restricted portfolio is also smaller. From the point of view of the fund manager, though, there will be a positive "shadow cost" to a binding risk constraint.

Such a cost can be measured in different ways, but perhaps the most useful measure is the reduction in the overall Sharpe ratio that would induce the fund manager to choose an unrestricted portfolio along the same indifference curve as the constrained one.[7] Graphically, this is given by the differences in slope between the actual and the "shadow" efficient frontiers in Figure 5.

As a numerical illustration, suppose that the investor's preferences are given by (1) with $\gamma = 1$, the net return on the safe asset is 0, and the returns on the primitive risky assets follow a multivariate t distribution with 8 degrees of freedom such that $\boldsymbol{\mu}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\iota} \neq 0$. If the maximum Sharpe ratio attainable is .5, which is roughly the value that reflects the performance of the US stock market over the period 1977 to 1997, an unconstrained investor will put 25% of her wealth in the optimal portfolio of risky assets, and the rest in the safe asset. As a result, the expected excess return on her optimally chosen portfolio will trivially be 25%, with a standard deviation of 50%. In contrast, an investor who must limit to 5% the probability of her losses exceeding 20%, will invest 91% of her initial wealth in the safe asset, a proportion that increases to just above 95% if the probability of incurring in such losses can only be 1%.[8] In the first case, the "shadow" Sharpe ratio is .384, while in the second case is .299, which implies that the shadow cost of a VaR constraint can indeed be substantial.

3. SUMMARY

In this paper I provide a simple unifying approach to static Mean-Variance analysis and Value at Risk, which can be used to carry out informed investment decisions within the well-known mean-variance allocation framework that satisfy the Value at Risk restrictions often imposed on fund managers by internal or external regulators.

This is achieved by introducing a new type of line to the usual mean - standard deviation diagram, which I call *IsoVaR*, which represents all the portfolios that share the same Value at Risk (as a proportion of the available funds) for a fixed probability level α. Given some pre-

---

[7] For instance, for the mean-variance preferences in (1), the "shadow" Share ratio, s_{p*} say, would be given by the formula $s_{p*} = \sqrt{2\gamma w_{r_p}^I [1 - (\gamma w_{r_p}^I)/(2s_p^2)]}$ = as long as $w_{r_p}^I < w_{r_p}^{MV}$.

[8] The corresponding pairs of expected excess returns and standard deviations (in percentage terms) will be (9, 18) and (5, 10) respectively.

Mean-Variance portfolio allocation...

specified risk control parameters, I show that the restricted portfolio allocation decision can be solved by simply superimposing the investors' mean-variance risk preferences on the section of the efficient frontier that is delimited by the appropriate *IsoVaR*. Therefore, the optimal mix of risky assets will be preserved, the only effect of the VaR constraint being a proportional reallocation of funds from the risky assets to the safe one.

Importantly, I show that from the point of view of the fund manager, there will be a "shadow cost" to a binding VaR constraint. This cost simply reflects the fact that the VaR restriction is imposed on the fund manager by an internal or external regulator, whose risk attitudes and objectives are different from the risk attitudes and objectives of the manager.

REFERENCES

Alexander, G.J. and A.M. Baptista. 2000. "Economic Implications of Using a Mean-VaR Model for Portfolio Selection: A Comparison with Mean-Variance Analysis", Working Paper, University of Minnesota (December).

Campbell, J.Y. and L.M. Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.

Chamberlain, G. 1983. "A Characterization of the Distributions that Imply Mean-Variance Utility Functions", *Journal of Economic Theory* vol. 29, no. 1 (January): 185-201.

Cuoco, D., He, H. and Issaenko, S. 2001. "Optimal Dynamic Trading Strategies with Risk Limits", Working Paper, Yale School of Management (May).

Dowd, K. (1999): "A Value at Risk Approach to Risk-Return Analysis", *Journal of Portfolio Management*, vol. 25, no. 4 (Summer): 60-67.

Fang, K.-T., S. Kotz, S. and K.-W. Ng. 1990. *Symmetric Multivariate and Related Distributions*, Chapman and Hall.

Gourieroux, C. and A. Monfort. 2001. "The Econometrics of Efficient Portfolios", Working Paper, CREST-INSEE (February).

Ingersoll, J.E. (1987): *Theory of Financial Decision Making*, Rowman and Littlefield, Maryland.

Jorion, P. (1997): *Value at Risk*, McGraw-Hill, New York.

Leibowitz, M.L. and S. Kogelman, S. 1991. "Asset Allocation under Shortfall Constraints", *Journal of Portfolio Management* vol. 17, no. 2 (Winter): 18-23.