September 2004

The generalised value at risk admissible set: constraint consistency and portfolio outcomes.

Roger J. Bowden[1]

Abstract

Generalised value at risk (GVaR) adds a conditional value at risk or censored mean lower bound to the standard value at risk and considers portfolio optimisation problems in the presence of both constraints. For normal distributions the censored mean is synonymous with the statistical hazard function, but this is not true for fat-tailed distributions. The latter turn out to imply much tighter bounds for the admissible portfolio set and indeed for the logistic, an upper bound for the portfolio variance that yields a simple portfolio choice rule. The choice theory in GVaR is in general not consistent with classic Von Neumann Morgenstern utility functions for money. A re-specification is suggested to make it so that gives a clearer picture of the economic role of the respective constraints. This can be used analytically to explore the choice of portfolio hedges.

Key words: Admissible set, censored mean, conditional value at risk, effective utility functions, generalised value at risk, hazard functions, hedging, portfolio choice, value at risk.

JEL numbers: G10,G11,D68,C61

---

[1] The author is with the School of Economics and Finance, the Victoria University of Wellington, P.O. Box 600 Wellington New Zealand, email roger.Bowden@vuw.ac.nz or roger.bowden@kiwicap.co.nz.
Special thanks go to Jennifer Zhu for some very useful background discussions and comments.

1 Introduction

Value at risk was originally conceived as a portfolio diagnostic, such that the probability of portfolio value falling below the designated value at risk corresponds to a chosen significance level (e.g. JPMorgan 1994 et seq., Butler 1998, Best 1999, Jorion 2002). Following its adoption by bank regulators in such contexts as the European directive and related Basle options, the methodology spread to other sectors such as funds or trustee managers and even non financial corporations. Objects of concern encompassed further domains like cash flow or earnings at risk, or even ‘drawdown at risk’ (referring to deposit outflows from fund managers).

A natural adaptation was to optimal portfolio design, with the idea that portfolio managers should maximise some welfare metric, e.g. the portfolio return, subject to a given value at risk. In addition to constraints or criteria based on the value at risk, some recent authors have proposed that an additional or alternative requirement be imposed, namely that the expected shortfall, given that the value at risk limit has been breached, should be bounded by a user-chosen number. The reasoning is that fat tailed return distributions can have a lot of sting remaining in the left hand tail, even if the probability of a return or value in that region is limited to 5%, say. The conditional value at risk (CVaR) is the expected portfolio value or return, given that it lies within the VaR critical region. By limiting this value, the manager can choose portfolios to minimise the damage remaining in the left hand tail. For contributions along these lines see e.g Andersson et al (1999) or Rockafellar and Uryasev (2000). In what follows, by ‘generalised value at risk’ is meant either or both VaR and CVaR, and any related methodologies (abbreviated GVaR, to coin yet another acronym!).

The general idea of setting VaR and CVaR as portfolio choice parameters certainly has appeal. The concepts are widely understood and therefore agreement can be obtained the more readily regarding what acceptable values might be. This would solve an under-recognised conceptual problem in risk management, namely that of securing common ground among the stakeholders, whether the latter are shareholders or financial regulators. One could indeed conceive of a risk management directive that simply says ‘maximise the expected return, subject to given bounds for the VaR and the CVaR’. In other words, the entire burden of risk aversion would be devolved to the generalised value at risk parameters. Once these are satisfied, the manager can effectively think in risk-neutral terms, ‘constrained risk neutrality’.

However there is a prior problem with risk management rules of this kind: do they work, and are they soundly based in the theory of choice under risk? Existing contributions have tended to concentrate on computing methods or particular applications, without going too deeply into the underlying conceptual issues. A problem is that the VaR and CVaR constraints, even though compatible in the programming sense of the existence of solutions, may nevertheless be economically inconsistent. The present paper is directed at a range of issues of the latter kind. A checklist of such concerns, together with some findings of the paper, is as follows:

(a) Conditions under which CVar and VaR might be mutually redundant. In optimisation terms, this would mean that likely solutions would see one constraint binding but not the other. It is shown that there is nothing logically inconsistent with both, so that a non empty feasible or legal set does exist, so that constraint qualification is not a problem. However, in most situations, one or other constraint is unlikely to be binding, and it is likely to be the VaR constraint. As it stands, it is the CVaR constraint that is the tough one: we shall suggest it is too demanding and needs reformulation.

(b) The paper explores the relationship of generalised value at risk to choice theory under risk, and in particular whether it is consistent with the Von Neumann – Morgenstern utility function for money, which forms the foundation of choice under risk. The effective utility of money is introduced, as a construct backed out of the saddle point property of the optimised Lagrangean, in the concave programming problem that underpins the generalised value at risk approach to portfolio selection. While the VaR constraint is perfectly consistent with this theory, CVaR is not. The problem arises because the two constraints are economically inconsistent if set independently of one another. They can be made consistent with a modification to the CVaR constraint. When this is done, one ends up with a nice separation of their respective roles. The GVaR-equivalent Von Neumann Morgenstern utility function can then be used to illuminate issues such as the optimal choice of portfolio hedge.

(c) What is the relation ship of generalised value at risk with more familiar ideas from statistics and reliability theory? VaR itself is simply a basic statistical construct derived from hypothesis testing and confidence levels. However, CVaR is a bit more than this, and the present paper links the idea of CVaR with hazard functions, implicitly exploring the use of the hazard function as a welfare indicator. This idea works for the normal distribution but not for more fat tailed distributions.

(d) Would a constrained risk neutral approach to risk management, along the lines earlier suggested, result in a unique solution or even any solution at all? In some situations, such as normally distributed security returns, it can. However, such a solution is likely to be characterised by a redundancy of the VaR constraint. It effectively amounts to maximising the portfolio variance along the mean variance efficient frontier, subject to the bound placed by the CVaR limit. The latter is in fact a fixed number for the logistic density of returns, which is leptokurtotic (fat tailed).

(e) A variety of other issues are examined, more or less as consequences of the preceding. One is the value of hedging. Hedge effectiveness is usually measured in terms of the conditional expectation of the hedge instrument return, given the return to be hedge. The associated empirical methodology is OLS or variants. However, once CVaR is introduced, hedge effectiveness become a matter also of the censored mean difference, which is identical with the ordered mean difference (Bowden, 2000). Hazard theory is a good framework for thinking about such problems. Another issue is whether GVaR can be reconciled with mean variance analysis. 'Delta-normal' VaR proceeds by first computing the means and covariance of factor exposures, so there is natural link. As indicated under item (c) above, the relationship can be developed further.

The scheme of the paper is as follows. Section II sets out the various forms of the GVaR portfolio optimisations. The CVaR constraint is identical with the censored mean evaluated at the VaR critical point, which enables a first look at portfolio choice. The resulting admissible set is explored for two representative cases in section III, namely the normal and logistic distributions, the latter chosen as a representative fat-tailed density. The binding nature of the constraints is analysed within this framework, as is the existence of bounds, including the logistic upper limit for the variance, and the relationship with statistical hazard theory. Section IV backs out the effective utility function and analyses compatibility with the Von Neumann-

Morgenstern theory of choice under risk. Section V suggests an effective utility function that overcomes some of the difficulties. It is also noted that that GVaR type problems necessitate an alternative empirical hedge methodology for identifying effective hedges and other enhancements. Section VI offers some conclusions.

2. GVaR decision problems

The generalised value at risk optimisation problems that have appeared in the literature have taken a number of forms, some mathematically dual to others, and with a variety of objective functions. The original formulations were usually cast in terms of portfolio values, as this is natural context for bank management. However, a reasonably embracing formulation could be written in terms of portfolio returns, *R*, and in what follows, we shall lose little generality by assuming that the objective function is simply *E[R]*. Alternative versions of the objective function involve minimising the value at risk or the conditional value at risk, and these will also be discussed.

2.1. The basic optimisation

The manager sets a lower critical point $R = r_L$ (usually negative) together with a significance level α and requires that the probability of a return less than r_L is no more than α. In other words, if *F (•)* denotes the distribution function of portfolio returns, we require $F(r_L) \le \alpha$. In addition, the CVaR constraint looks at the conditional expectation $E[R|R \le r_L]$ and requires that this be at least a given value *v*, say. The problem is to choose decision parameters **z** (e.g. portfolio proportions) to maximise the objective function subject to the GVaR constraints. Thus the optimisation problem can be written as follows:

Given a preset value at risk r_L , a significance level α, and a return number $v < r_L$:

$$\max_z \; E[R]$$
$$\textit{subject to} \qquad (1)$$
$$(i) \quad F_R(r_L) \le \alpha$$
$$(ii) \quad E[R|R \le r_L] \ge v$$

and possibly other constraints (e.g. proportions add up to 1). The latter can remain implicit in what follows, though the issue is revisited in section III. To avoid proliferation of subscripts, $F_R(r)$ will often be written just as F(r), when the context is clear. The corresponding density, assuming it exists, is written as $f_R(r)$ or just $f(r)$.

Alternative decision problems

Other versions might proceed to minimise the value at risk subject to achieving a minimal expected return, or alternatively maximise the conditional value at risk. For instance, one could have

$$
\text{(2a)} \quad
\begin{aligned}
&\min_{z, r_L} \; r_L \\
&\textit{subject to} \\
&(i) \quad F_R(r_L) = \alpha \\
&(ii) \quad E[R] \geq m \\
&(iii) \quad E[R \mid R \leq r_L] \geq v
\end{aligned}
$$

where *m* is some pre-set 'satisficing' level of expected return. Or the objective could be cast as maximising the conditional value at risk:

$$
\text{(2b)} \quad
\begin{aligned}
&\max_{z} \; E[R \mid R \leq r_L] \\
&\textit{subject to} \\
&(i) \quad F_R(r_L) \leq \alpha \\
&(ii) \quad E[R] \geq m
\end{aligned}
$$

In certain cases, the minimisation problem (2a) may emerge as a dual (Rockafellar (1968)) to the primal (1). In general, however, this is not true for more or less arbitrary user-defined values of the constants *m* or *v*, and duality will not be assumed in what follows.

The original version (1) corresponds to the way that most asset portfolio users would define the decision context. The basic objective of the firm is to maximise portfolio value. Constraints like VaR and CVaR are seen as administrative, regulatory, or 'legal' constraints, and *admissible* portfolios are those that satisfy them. Subject to these constraints, the manager seeks to maximise some welfare function that for present purposes will be taken as the expected return. Much the same ideas will apply even if the manager is risk averse, so that instead of maximising the expected return *E[R],* he or she wishes to maximises the expectation of some concave utility function *E[U(R)]*. However, there is independent virtue in simply assuming expected return as the utility function, because we can check out an idea earlier mentioned, namely that the burden of risk management could be entirely set via the GVaR constraints, leaving the agent free (within the admissible set) to otherwise act as risk neutral.

2.2. The CVaR as the censored mean

The function

(3) $\gamma(r) = E[R|R \leq r]$

is the censored mean of the distribution of R. Equivalently, it is the unconditional mean of the truncated distribution whose density is given by

$$f_c(R) = \frac{f(R)}{F(r)}; \quad R \leq r .$$

Thus

(4) $$\gamma(r) = \frac{1}{F(r)} \int_{-\infty}^{r} Rf(R)dR .$$

For a pre-set number $r = r_L$, long left hand tail distributions such that the density declines only slowly to the left of $R = r_L$ will have a numerically larger value of $\gamma(r_L)$ and hence be more damaging.

An equivalent formulation runs in terms of the cumulated distribution function familiar from second order stochastic dominance theory, namely:

$$\Phi(r\ ) = \int_{-\infty}^{r} F(R)dR .$$

For suitably regular distributions, we have

(5) $$\gamma(r) = r - \frac{\Phi(r)}{F(r)} .$$

In particular, if $F(r_L) = \alpha$ then comparing two alternative return distributions with the same value at risk, A will have larger value of the censored mean than B if $\Phi_A(r_L) < \Phi_B(r_L)$. Thus if B stochastically dominated A, then the hazard associated with B would be less for any value at risk. However in what follows we do not assume stochastic dominance or work further in such terms.

The function $\gamma(r)$ is a monotonically rising function of r tending asymptotically to the unconditional mean E[R] = μ_R , and tangential to the 45% degree line as R→ - ∞ . It is sketched in figure 1 as globally concave, which is usually the case.

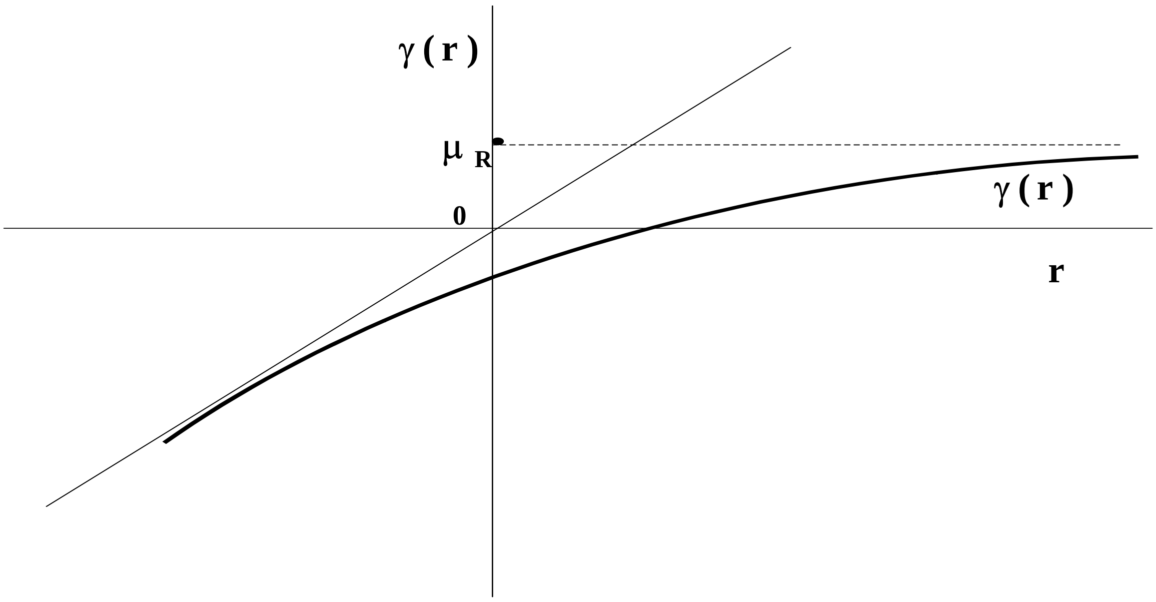

Figure 1: The censored mean function

2.3. How the constraints work

Figure 2 illustrates the censored mean functions for three alternative portfolios z_A , z_B , z_C generating portfolio returns R_A , R_B , R_C , respectively. Portfolio C has the highest mean return, marked in as $\mu_C = E[R_C]$. But its conditional or censored mean diminishes very quickly for low return values. In particular, the value of $\gamma_C(r)$ at the VaR point $r = r_L$ is less than the preset magnitude v, so it is not ‘legal’. Portfolios A and B are legal, in this sense, but A is to be preferred of the two, as its unconditional mean is higher. Appendix I is a corresponding illustration for the ‘dual’ version problem (2a) above.

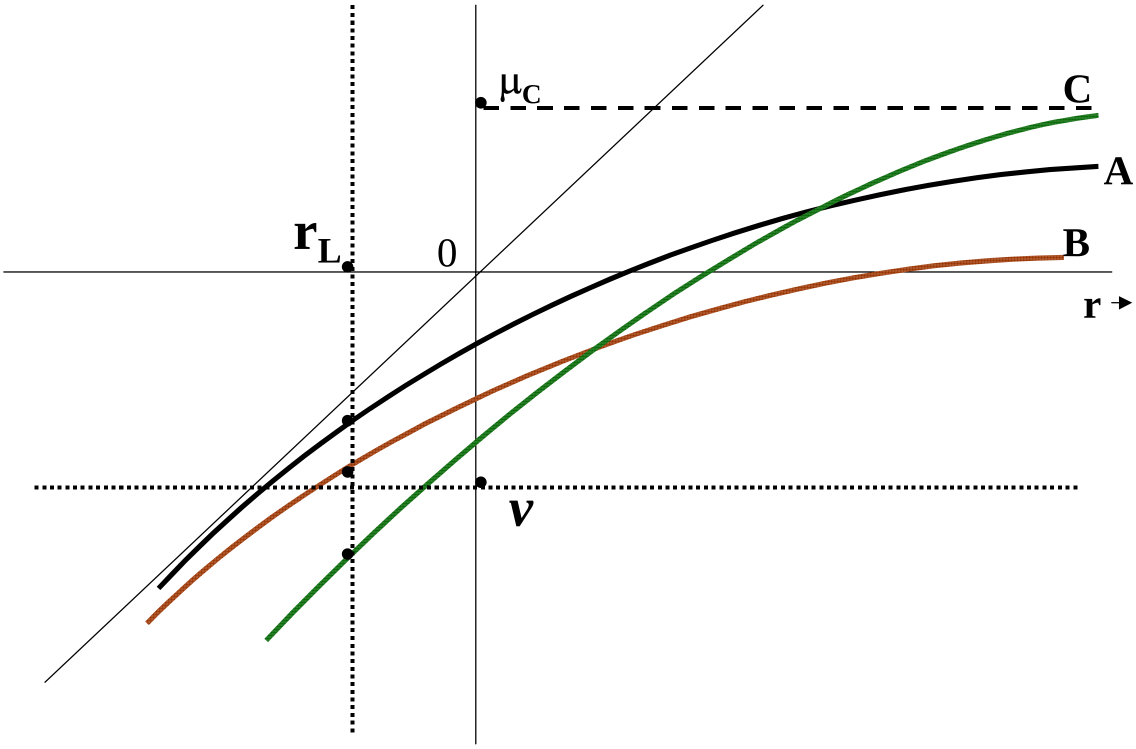

Figure 2: Portfolio comparison

3. The admissible set and its implications

The admissible set, or 'legal set', will refer to the portfolios that satisfy both the VaR and CVaR constraints. It is not quite the same as the feasible set, which refers to the portfolio proportions ***z*** leading to admissible portfolios and in addition satisfying any further constraints e.g. $\sum_i z_i = 1$. The purpose of this section is to explore general considerations that might serve to bound the admissible set, how this might depend upon the nature of the returns distribution, the relationship with hazard theory; and further issues such as whether the CVaR and VaR constraints can both be binding at the same time.

For this purpose, it is convenient to use two-parameter distributions, such as the normal and logistic, choosing the parameters as the mean μ_R and standard deviation, σ_R . Both these distributions admit a standardisation of the form

(6) $$Y = \frac{R - \mu_R}{\beta},$$

where β is a variance- related constant. For the normal distribution, $\beta = \sigma_R$ and the distribution of Y has mean zero and variance unity. The following lemma will be useful in what follows.

Lemma 1

Let γ_S refer to the censored mean or hazard function of the standardised distribution. If $y = \frac{r - \mu_R}{\beta}$ *, then in the original units, the censored mean is given by*

(7) $$\gamma(r) = \mu_R + \beta\gamma_S(y)$$

[Proof: Follows from the definition (4) and transformation of variables].

3.1. The normal distribution

The following result summarises the properties of interest for the current discussion.

Proposition 1:

Suppose R is N(μ_R , σ_R^2) , let $Y = \frac{R - \mu_R}{\sigma_R}$, *and write the standard normal distribution and density as N(Y), n(Y) respectively. Then the censored mean at R = r is given by* $\gamma(r) = \mu_R + \sigma_R\gamma_S(y)$

where $y = \frac{r - \mu_R}{\sigma_R}$ *and* $\gamma_S(y) = -\frac{n(y)}{N(y)}$.

Proof:

Consider first the censored mean for the standard normal distribution, which is given by

$$\gamma_S(y) = \frac{1}{N(y)}\int_{-\infty}^{y} Yn(Y)dY .$$

As $n'(Y) = -Yn(Y)$, it follows that $\gamma_S(y) = -\frac{n(y)}{N(y)}$.

The desired result follows from Lemma 1.

□

Given a VaR point r_L , let $y_L = \dfrac{r_L - \mu_R}{\sigma_R}$ be the corresponding standardisation. Also let y_L^{α} be the unit normal lower one tailed critical point for significance level α, e.g. (–1.65) for α =5%.

One can now proceed to locate admissible regions in the (μ_R , σ_R) plane, with σ_R taken along the horizontal axis. Using conditions (1)(i), (ii) and Proposition 1, the admissable regions taken for each constraint separately are located as follows:

(8a) *VaR*: $\mu_R \geq r_R - \sigma_R y_L^{\alpha}$

(8b) CVaR: $\mu_R \geq v + \sigma_R \dfrac{n(y_L)}{N(y_L)}; \quad y_L = \dfrac{r_L - \mu_R}{\sigma_R}$

It will be noted that the CVaR region is implicit in μ_R . Its boundary can be found by replacing the above equality with an equality and solving for μ_R in terms of σ_R. It is a convex rising function of σ_R.

Figure 3 locates the feasible region for the stated parameters as the intersection of the two feasible regions for each constraint, for the given user parameters.

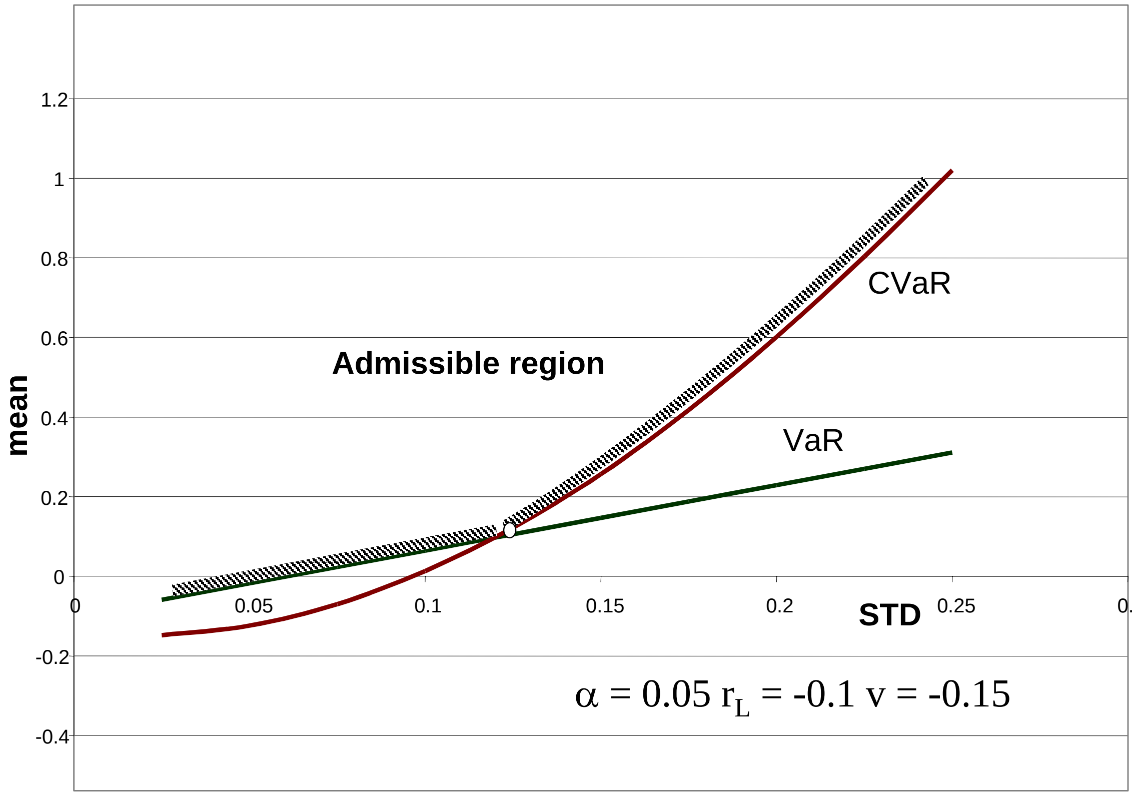

Figure 3: Admissible region for the normal distribution

Figure 4 combines this with the locus of (μ_R ,σ_R) combinations that might be possible with different portfolio proportions. The hatched locus corresponds to the efficient frontier in standard mean-variance analysis. Given the objective to maximise the mean, the optimal GVaR portfolio point is indicated as the point A. This corresponds to maximising the STD of the portfolio subject to the CVaR constraint. In this example the VaR constraint will not be binding at the optimum.

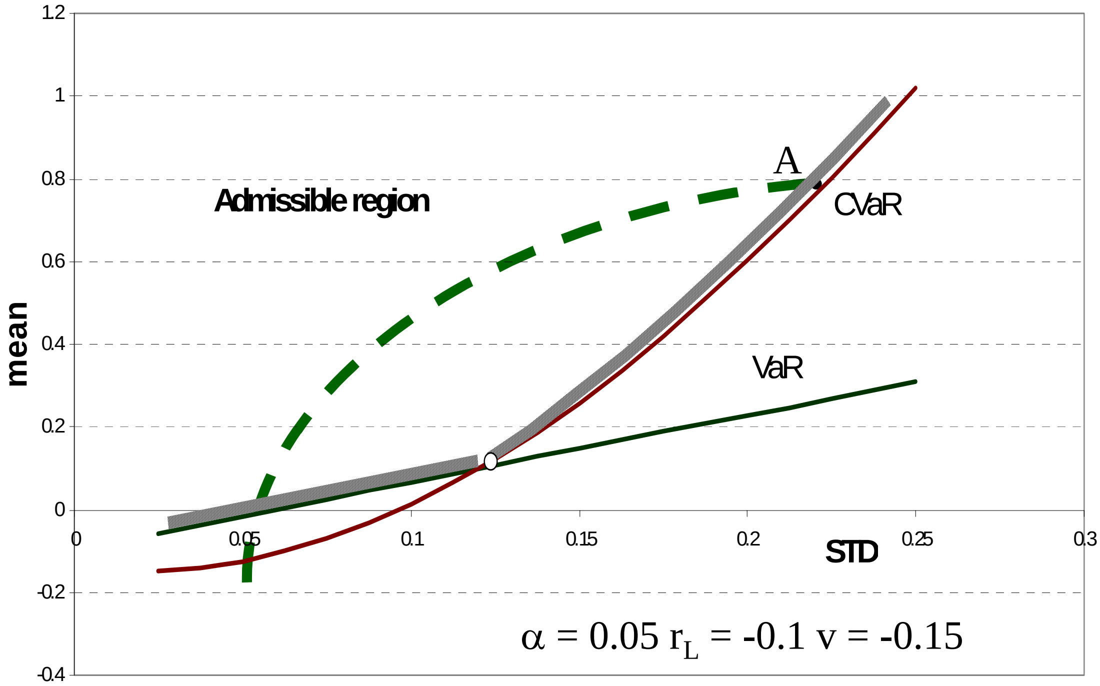

Figure 4: Combining with mean variance analysis

3.2. Relationship with hazard theory

The function $H_n(y) = \dfrac{n(y)}{1 - N(y)}$ is the normal hazard function of statistical survival and reliability theory (e.g. Mann *et al* 1974). Hence in proposition 1, the function $h_n(y) = \dfrac{n(y)}{N(y)} = H_n(-y)$, which can be regarded as the hazard function for (-*y*), i.e. going in the reverse direction. Once could perhaps call h_n the 'reverse hazard function'. Thus for the normal distribution, the censored mean is the negative of the reverse hazard function. One could imagine that the probability of the institution going bankrupt increases as y diminishes, or the hazard or marginal probability of 'death' associated with any given value increases. Figure 5 plots the standard normal hazard function in the required reverse form $h_n(y)$. In this sense, CVaR can be said to capture the 'hazard' of the situation, at least for the normal distribution. But as we shall later see, there are severe limitations to the implicit identification of the hazard function with economic welfare for more general distributions.

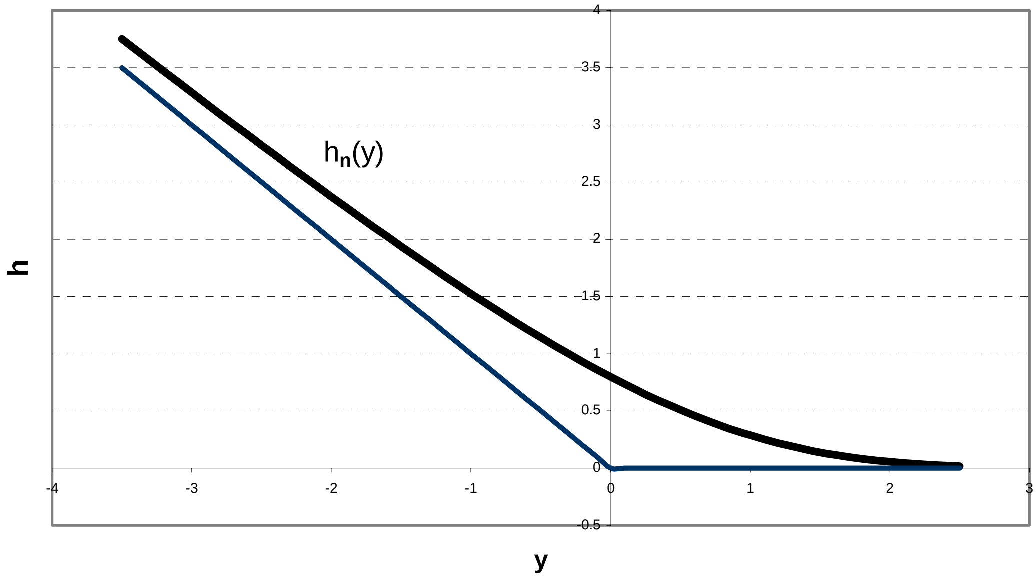

Figure 5: Standard normal reverse hazard function

The normal distribution has the convenient property of being preserved (or 'stable') under linear combinations. This meant that in figure 4 we were able to consider the admissible region in conjunction with the mean-variance set traced out by alternative portfolios. The next distribution, namely the logistic, does not have the preservation property, so the admissible set is not a global one for all possible portfolios. However, it does reveal what happens for fatter tail distributions, namely that the variance bound becomes extreme.

3.3. The logistic distribution

The logistic distribution function is defined by

$$F(R) = \frac{1}{1 + \exp[-\frac{(R - \mu_R)}{\beta}]} ; \sigma_R = \beta\pi / \sqrt{3} .$$

A convenient standardisation is

$$Y = \frac{R - \mu_R}{\beta} \text{ with } F_s(Y) = \frac{1}{1 + e^{-Y}}.$$

The logistic density is symmetric, but has fatter tails than the normal: a kurtosis coefficient of 4.2, versus 3 (Johnson and Kotz 1970). This makes it useful for present purposes, as CVaR has in effect been designed for such contingencies.

For the logistic distribution, the censored mean function at R = r is given by

(9) $$\gamma(r) = r + \frac{\beta}{F\ (r)} \log[1 - F(r)].$$

Appendix II proves the above result and also establishes the admissible set, which is depicted in figure 6.

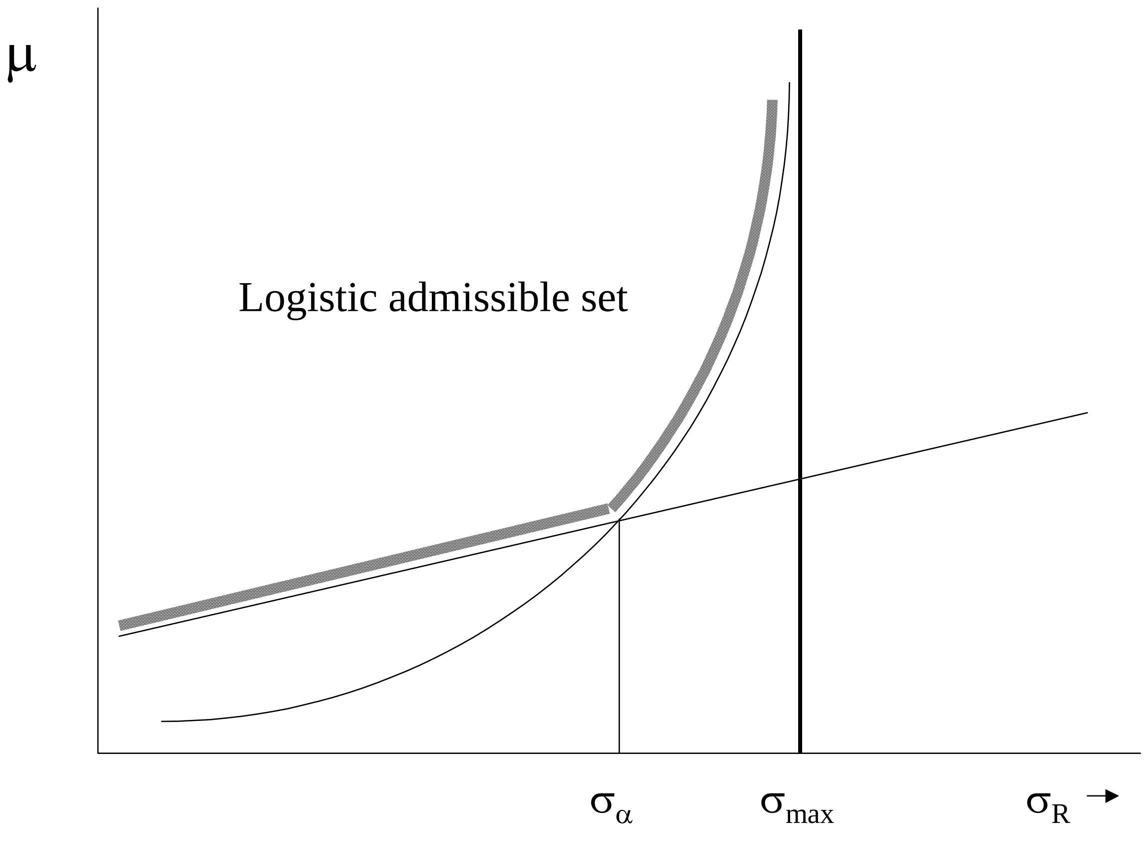

Figure 6: Existence of the logistic admissible set

Perhaps the most interesting feature is the existence of a firm upper bound for the variance, given by:

(10) $$\sigma_R \le \frac{\pi(r_L - v)}{\sqrt{3}} \approx 1.814(r_L - v).$$

Thus for $r_L = -0.10, v = -0.15$, the limiting STD = 0.09.

Comparing this with the same settings for the normal distribution as in figure 3, the boundary effect is very restrictive. Figure 6 suggests that the portfolio solution may have quite a simple first approximation: fix the STD at the logistic upper bound $\sigma_{\max} = 1.814(r_L - v)$ and maximise the portfolio mean. The resulting portfolio may not be admissible, but it might well be close.

It is a reasonable conjecture that similar properties hold for other fat-tailed distributions. If there is any reason to suspect that security returns are not going to be normally distributed, it may be quite unsafe to adopt high variance portfolios even though they might satisfy the admissible set for the normal distribution. A conservative approach might simply be to adopt the logistic upper bound as representative of other fat-tailed distributions, and limit the portfolios to those that satisfy the logistic STD bound.

The logistic hazard function

The hazard function for the standardised logistic is

$$H_S(y) = \frac{f_S(y)}{1 - F_S(y)} = F_S(y),$$

so that the reverse hazard function is $h_S(y) = \dfrac{f_S(y)}{F_S(y)} = 1 - F_S(y)$. Unlike the normal distribution, this is not equal to the negative of the censored mean. The relationship is

$$\gamma_S(y) = y + \frac{\log(h_S(y))}{1 - h_S(y)}.$$

Figure 7 compares $h_s(y)$ with $-\gamma_S(y)$ for the logistic. Evidently the censored mean is a much more demanding loss function than is the hazard function. We can now see the general problem with the hazard function as a measure of welfare loss. Because of the long left hand tail, the logistic 'survival' always remains significant, whereas with a shorter tail, the normal death probability becomes overwhelming, and the hazard correspondingly explodes. The problem is that economic survival does not coincide with statistical survival. Thus the bank or fund whose returns follow a logistic distribution may be bankrupt long before it is flagged as a concern by the statistical hazard rate. This is why the censored mean is a better indication of potential loss than the statistical hazard function.

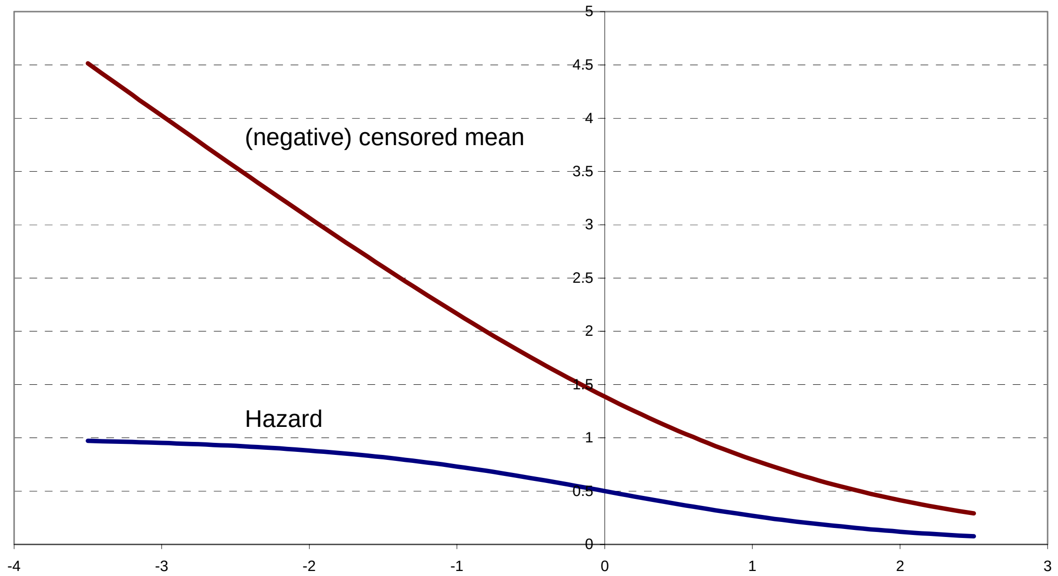

Figure 7: Logistic reverse hazard and censored mean compared

It can be concluded from the above that the CVaR constraint is much more demanding in the case of the logistic distribution. The VaR constraint is likely to be redundant at the optimum, and there is also a fixed upper bound on the portfolio variance.

4. Effective utility functions

Optimisation problems such as those in mean variance analysis can always be regarded as maximising an expected utility function subject to any portfolio constraints. The effective utility function (up to inessential additive and slope constants) is of the form $U(R) = R - cR^2$ for some constant c. This well known form can be backed out of the standard Lagrangean for the optimisation problem using the methods described below. It is of interest to identify the corresponding effective utility function for the GVaR optimisation problem.

The GVaR decision problem (1) (section 2) may be written in the equivalent form:

$$
(11) \quad \begin{aligned} & \max_z \; E[R] \\ & \textit{subject to} \\ & (i) \quad F_R(r_L) \le \alpha \\ & (ii) \quad E[R | R \le r_L] F_R(r_L) \ge v F_R(r_L) \end{aligned}
$$

Multiplying constraint (ii) in version (1) by $F_R(r_L)$ to get version (11) will make no difference to the solution for the optimal portfolio proportions. The associated Lagrange multipliers will have to be reinterpreted, but will remain unchanged in sign. In addition we may, if so wished, assume portfolio proportion constraints such as $\sum_i z_i = 1$ or $z_i \ge 0$.

The Kuhn-Tucker Lagrangean for the optimisation problem (11) may be written as:

(12)

$$
L(z;\mu,\lambda) = E[R] - \mu F_R(r_L)(v - E[R | R \le r_L]) + \lambda(\alpha - F_R(r_L)) \quad [+\rho(\sum_i z_i - 1) + \sum_i \theta_i z_i]
$$

where the further terms in square brackets may be added or deleted as appropriate. The Lagrange multipliers λ, μ are those of primary interest here, so we suppress the others (ρ, θ). At the optimum solution, we must have:

$$
(13) \quad \begin{aligned} & \mu \ge 0; \quad \mu F_R(r_L)(v - E[R | R \le r_L]) = 0 \\ & \lambda \ge 0; \quad \lambda(\alpha - F_R(r_L)) = 0 \end{aligned}
$$

the complementary slackness conditions. According to the saddle point property, the optimum solution may be regarded as solving

$$
\max_z \min_{\mu,\lambda} L_a(z;\mu,\lambda),
$$

subject to any further constraints on the z_i .

The effective utility function

Introduce the unit step function

$$
\begin{aligned} SF(x) \; & = 1 \Leftrightarrow x > 0 \\ & = 0.5 \Leftrightarrow x = 0 \\ & = 0 \Leftrightarrow x < 0. \end{aligned}
$$

It follows that

$$
(14) \quad E[SF(r_L - R)] = F_R(r_L).
$$

Also from equation (4) section 2, we have

(15) $$E[R|R \le r_L] = \frac{E[R\,SF(r_L - R)]}{E[SF(r_L - R)]} = \frac{E[R\,SF(r_L - R)]}{F_R(r_L)}.$$

Substituting (14) and (15) into (12), the Lagrangean can be written

$$L(z;\mu,\lambda) = E[R - \mu(v - R)SF(r_L - R) + \lambda(\alpha - SF(r_L - R))]$$
$$[+\rho(\sum_i z_i - 1) + \sum_i \theta_i z_i].$$

Thus the original decision problem (1) is effectively the same as maximising the expected value *E[U(R)]* of a utility function *U(R)* which can itself be written

(16) $$U(R) = R - r_L - \mu(v - R)SF(r_L - R) + \lambda(\alpha - SF(r_L - R)),$$

where the additional constant r_L has been introduced for later convenience. In this maximisation it can be assumed that the Lagrange multipliers μ, λ have been fixed at their optimal values. If this is the case, then maximising the expected value of the effective utility function (16) will yield the same solution as the full decision problems (1) or (11).

4.1 Analysis of the effective utility function

The zonal behaviour of *U (R)* as a function of *R* is as follows:

(a) $R > r_L$

(17a) $U(R) = R - r_L + \alpha\lambda$;

(b) $R = r_L$

(17b) $U(R) = \alpha\lambda - \frac{1}{2}\pi$;

(17c) $R < r_L$

(17c) $U(R) = (1+\mu)(R - r_L) + \alpha\lambda - \pi$,

where $\pi = \lambda - \mu(r_L - v)$.

Figure 8a sketches the resulting function on the assumption that π > 0, i.e. that $\mu(r_L - v) < \lambda$ .The effective utility function is basically equivalent to R above the VaR point r_L . At that point it drops down sharply, and as R falls further, the function regresses at a faster rate than R. The drop $\pi = \lambda - \mu(r_L - v)$ at the VaR could be interpreted as a lump sum penalty where the VaR constraint is breached. The rate at which *U(R)* falls thereafter depends upon the dual Lagrange multiplier μ, and the slope *(1+μ)* is numerically larger than in the upper branch. Effectively this is telling the user to penalise return distributions that have a long left hand tail, just as one would expect from the CVaR constraint.

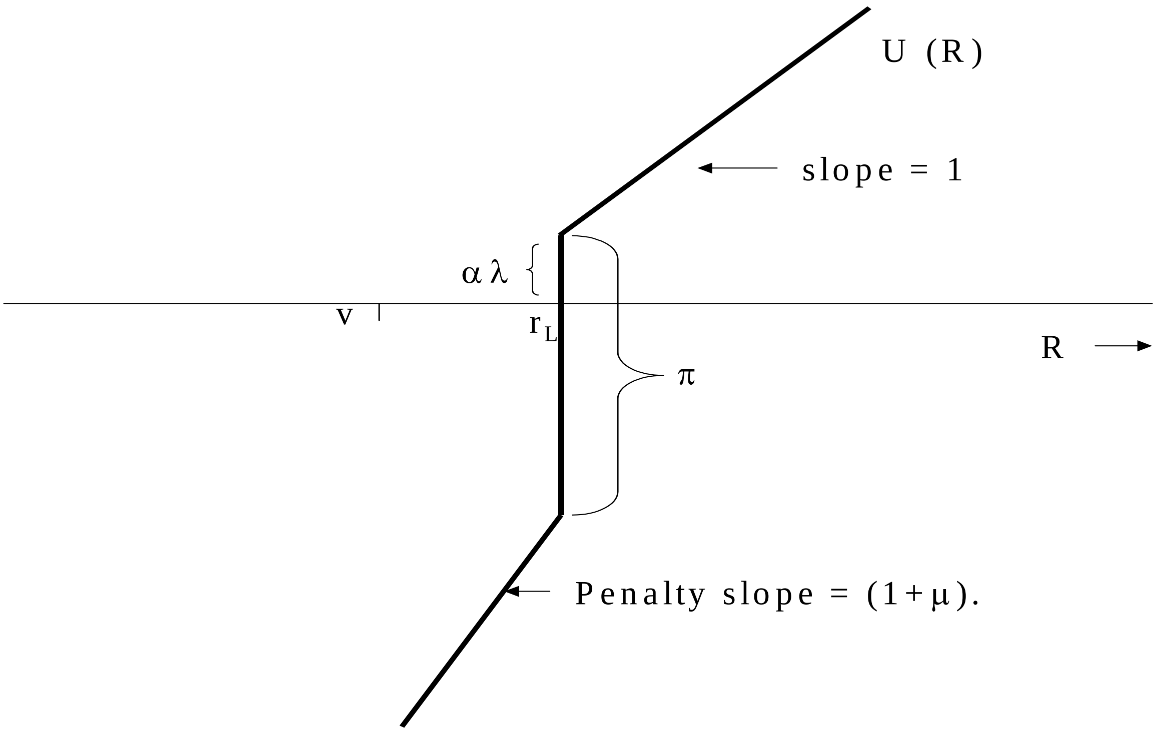

Figure 8a: The effective utility function with π > 0

However, one can also have π < 0, i.e. $\mu(r_L - v) > \lambda$. Figure 8b sketches the resulting effective utility function. In this case perverse behaviour appears just to the left of the VaR point, according to which the effective utility of money actually increases before resuming its downward path, at the penalty slope 1+μ. In this case, an equivalent Von Neumann-Morgenstern utility function does not exist.

The various possibilities are as follows:

(a) VaR is alone binding at the optimum. In this case π = αλ > 0. The effect is that of a pure lump sum penalty once the VaR point is reached, and there is no further slope penalty.

(b) CVaR is alone binding at the optimum. In this case $\pi = -\mu(r_L\text{-}v) < 0$. This is the perverse behaviour noted as for figure 6b.

(c) Both VaR and CVaR are binding at the optimum. In this case π could be of either sign and a proper Von Neumann- Morgenstern utility function may or may not exist.

Possibility (b) identifies the source of the trouble with the CVaR constraint. In effect, it is telling us that the latter is not properly integrated with the VaR constraint, so that one should not really be setting the parameter v, α independently of one another. As earlier remarked, case (b) is likely to hold for fat-tailed distributions. This means that the implied optimisation cannot be easily reconciled with the Von Neumann Morgenstern choice theory under risk.

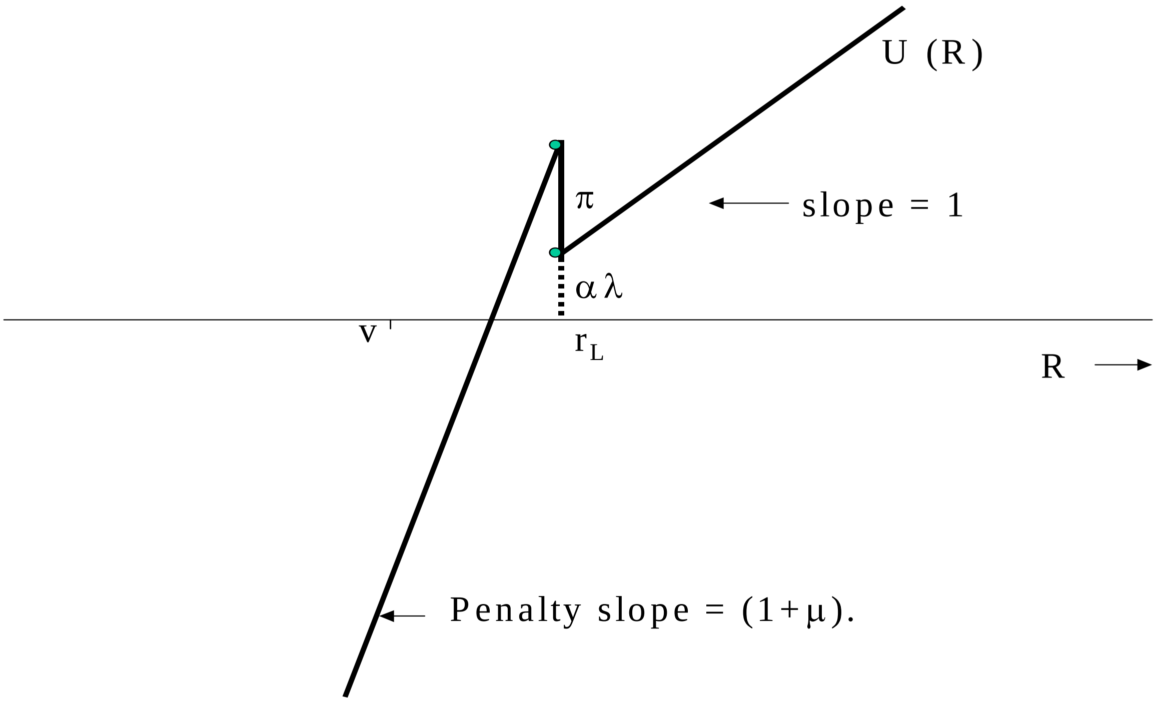

Figure 8b: The effective utility function with π < 0

4.2 Other versions of the decision problem

It is of interest to note that the version (2c) of the section II optimisation problem, which seeks to maximise the conditional value at risk, can result in a sensible Von Neumann - Morgenstern utility function. The decision problem may be written:

$$\begin{aligned}
&\max_z E[R|R \le r_L] \\
\text{(2a)}\quad &\textit{subject to} \\
&(i) \quad F_R(r_L) \le \alpha\,; \quad (\lambda) \\
&(ii) \quad E[R] \ge m\,; \quad (\theta)
\end{aligned}$$

where we have indicated relevant semipositive Lagrange multipliers, alongside each constraint. Following through the same procedure as above, the equivalent utility function may be written

$$U(R) = [\theta + SF(r_L - R)](R - r_L) - (\lambda - r_L)SF(r_L - R) + \lambda\alpha - (m - r_L)\theta\,.$$

Features are:

(a) The step down $\pi = \lambda - r_L$ at the VaR point $R = r_L$ is likely to have the right sign, especially if $r_L < 0$, as one should normally expect.

(b) To the left of the VaR point $R = r_L$, the slope *is* $(1+\theta)$ compared with θ to the right of this point, again indicating a slope penalty.

In this version, the rewards for exceeding the VaR point will be non zero only if the mean constraint (ii) is binding, i.e. if $E[R]=m$, the required minimum, in the optimal solution. This can therefore be regarded as a defensive portfolio strategy. There is no primary utility interest in exceeding the required minimum portfolio mean. This version does not correspond to the usual portfolio objective, which is to maximise some functional of returns, subject to administrative or regularity constraints.

5. GVaR-mimicking utility functions

In view of the problematic features noted above, one could initially adopt the more informal approach of a user-defined Von Neumann- Morgenstern utility function that captures the more essential features of the GVaR penalties. Once this form is decided, one could identify a set of constraints that captured correctly the desired behaviour.

Suppose that instead of the original version with user-defined numerical constraints (e.g. v), the manager elected to pre- set the penalty elements according to some preferred balance between the penalty for reaching the VaR limit and that for exceeding it (i.e. the CVaR element). The implicit utility function would then become explicit, of the form:

(18) $$U(R) = [1 + \mu SF(r_L - R)](R - r_L) - \pi SR(r_L - R)\,.$$

In expression (18), the user would set a VaR- type lump sum penalty π and the CVaR-type slope penalty $1+\mu$. The precise calibration depends on issues such as the existence of a regulatory penalty (higher value of π) or apprehensions about the potential length of the left hand distributional tail (higher μ). Figure 9 illustrates.

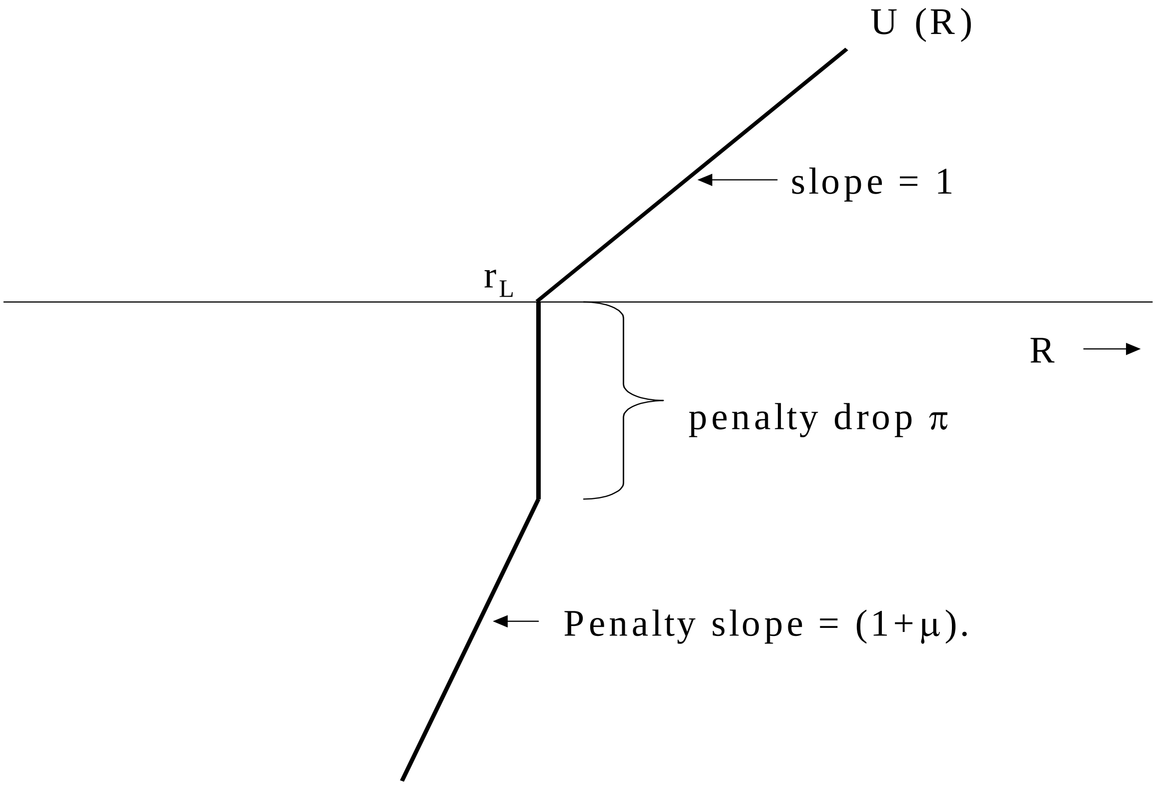

Figure 9: GVaR - mimicking utility function

5.1 Conformation with programming versions

The form (18) can be made to conform with mathematical programming versions by means of an alteration that points up the essential problems with the original. The GVaR mimicking utility function (18) can trivially be recast as

$$U(R) = R - r_L + \mu[SF(r_L - R)(R - r_L) + \delta] - \pi[SF(r_L - R) - \alpha],$$

where δ and α are positive constants, the latter corresponding to the VaR significance level. Maximising *E[U(R)]* would then correspond to the following programming problem:

$$\max_z \; E[R]$$

subject to

(19) $(i) \quad F_R(r_L) \le \alpha$

$$(ii) \quad E[R \mid R \le r_L] \ge r_L - \frac{\delta}{F_R(r_L)}$$

Referring back to the original problem (1), the CVaR constraint *v* is no longer a constant but is variable, of the form

$$v = r_L - \frac{\delta}{F_R(r_L)} .$$

In the revised version (19ii), the conditional expectation requirement is relaxed if the significance level attained at r_L is smaller. This is telling us that in the original version, keeping the CVaR constraint v as a constant is simply too tough a requirement for the inherent probability of the situation. A candidate solution that attains a low probability of being in the critical VaR region should not be penalised with an unreasonably high CVaR setting.

Or to put it in a nutshell, the VaR and CVaR settings should not be imposed in isolation from one another. Imposing the CVaR constraint in the form (19ii) leads to a much clearer separation of the roles of the two constraints.

(i) The shadow price (expected optimal return increment) of the value at risk constraint leads to the lump sum penalty π for violating it;

(ii) The shadow price of the revised CVaR constraint leads to the slope penalty once the VaR constraint has been violated.

In what follows, it is assumed that the economically consistent form of the CVaR constraint is utilised. This means that the effective utility function is of the GVaR-mimicking form (18).

5.2 The marginal value of hedging

Generalised value at risk attitudes towards risk imply that traditional empirical hedge theory needs to be supplemented with a cumulative element, reflecting the important placed on the left hand tail area. GVaR-mimicking utility functions provide a simple framework for evaluating a proposed portfolio hedge or enhancement in such terms. The same framework can cover issues of portfolio choice, e.g. whether a new security or asset will add any additional value to the portfolio.

Let R_b be the return on a base or benchmark portfolio and R_a denote the return on a proposed portfolio enhancement or hedge. The latter may or may not call for capital to be allocated. Forwards or swaps typically do not, while options do. If there is capital to be allocated, then the appropriate return contribution is the difference $R_a - R_b$ meaning that one funds the proposed addition by down-weighting (going short in) R_b . In the zero capital framework, replace the difference by just R_a; for a further discussion see Bowden (2003).

Assume provisionally that capital is to be allocated. In that case, the marginal value of the enhancement is given by the equivalent variation

(20) $$\tau_{a/b} = \frac{E[(R_a - R_b)U'(R_b)]}{E[U'(R_b)]} \ .$$

The metric τ forms the basis of ordered mean difference portfolio technology (e.g. Bowden (2000, 2005). It has a returns dimension, and measures the amount that the portfolio manager would have to be compensated to give up the opportunity of being able to invest a dollar in asset of return R_a. The numerator could also be interpreted in terms of the directional derivative:

$$\left.\frac{\partial E[U(R_b + \theta(R_a - R_b))]}{\partial \theta}\right]_{\theta=0}$$

The denominator normalises the directional derivative so that it assumes the desired returns dimension.

If asset R_a is to add anything to the base portfolio, then $\tau_{a/b}$ is non zero; positive means to go long or longer, and negative, go short. It is zero if asset R_a is already an optimised part of the base portfolio, so $\tau_{a/b} = 0$ indicates a portfolio equilibrium. The following result gives an explicit formula for the equivalent variation- it is interpreted below.

Proposition 2

Define the reverse hazard function[2] $h_b(r) = \frac{f_b(r)}{F_b(r)}$ *where* $F_b(r)$ *is the distribution function of* R_b *evaluated at* $R_b = r$ *and* $f_b(r)$ *is the corresponding density function evaluated at the same point. Then*

(21)

$$\tau_{a/b} = \frac{E[R_a - R_b] + F_b(r)\{\mu E[(R_a - R_b)|R_b \le r_L] + \pi E[(R_a - r_L)|R_b = r_L]h_b(r_L)\}}{1 + F_b(r_L)[\mu + \pi h_b(r_L)]}$$

Interpretation

The numerator of (21) contains all the important elements for assessing the effectiveness of the proposed enhancement:

(i) The first term is simply the mean difference relative to the existing portfolio.

---

[2] Technically this will be the reverse hazard function only if the density is symmetric. However it is useful to refer to it in such terms, even where this is not the case.

(ii) The second term, namely

$E[(R_a - R_b) | R_b \leq r_L]$

is the mean difference censored by the base portfolio value at its VaR point. The sample version is the ordered mean difference of R_a with respect to R_b, evaluated at the VaR point r_L .

(iii) The third term,

$E[(R_a - r_L) | R_b = r_L] h(r_b)$

represents the improvement precisely at the VaR critical point, weighted by the reverse hazard function for the base portfolio evaluated at that point. This is computed as the fitted value in the regression of $R_a - r_L$ on R_b, evaluated at the VaR point.

Terms (ii) and (iii) are themselves weighted by the probability that the benchmark return falls into the critical area $\{R_b \leq r_L\}$. The second term is of special interest because traditional least squares hedging theory deals only with the conditional expectation, in present terms an item like (iii). If there is a CVaR element involved then the truncated or censored mean becomes of relevance, taken over the whole critical area. Empirically, this can be computed using the ordered mean difference sample estimator. Other things being equal, one should hedge with a security whose OMD over the current benchmark is as large as possible in the critical region.

The preceding analysis is easily adapted to the case where R_a uses no capital, as in a forward purchase or sale. The corresponding equivalent margin is given by

$$\tau_{a/b} = \frac{E[R_a] + F_b(r)\{\mu E[(R_a) | R_b \leq r_L] + \pi E[(R_a) | R_b = r_L] h(r_L)\}}{1 + F_b(r_L)[\mu + \pi h(r_L)]} .$$

In this version, R_a is a pure zero-capital enhancement, but similar terms and interpretations apply.

6. Conclusions

By linking the generalised value at risk constraints to statistical constructs such as the censored mean and hazard functions, some of the issues raised in section 1 can be explored.

In general, both the VaR and CVaR constraints play a role, so that there is no automatic dominance of one by the other, over all types of distribution. For fat- tailed situations for which conditional value at risk is considered appropriate, it is likely that the value at risk constraint as such is likely to be non binding. This does not mean that the VaR constraint is strictly redundant, as it may nevertheless remain as a useful guide for iterative solution paths. However, another useful feature is the existence of a firm upper limit for the portfolio variance in such cases. This was derived in the case of a logistic distribution, but it seems a reasonable conjecture that a similar property holds for other fat tailed distributions. Indeed, one could think of a 'near enough' procedure of simply setting a variance bound at the logistic upper limit and proceeding to maximising the mean. It would be of interest to explore the robustness of this simple rule, compared with a full GVaR maximisation procedure.

Hazard functions fall short as a candidate for possible loss functions or constraints. For normal distributions they are identical with the CVaR constraint. This is not true for fat tailed distributions, wherein statistical hazard severely understates the true economic hazard.

Taken together, the above observations indicate that to assume normal distributions (as in the delta normal VaR methodology) may be dangerously misleading as to the binding status of constraints and their economic effect. This means that in any situation where fat tailed distributions are envisaged, an adequate GVaR analysis will require much more attention to distributional properties and implications. A simple check might be to compute the logistic upper limit and compare this with the variance that actually emerges from a delta normal based portfolio optimisation.

The existence is noted of determinate GVAR-constrained portfolio solutions even where the objective is simply to maximise the expected return. This suggests that GVaR optimisation might enable the burden or risk management to be devolved, as it were, to the GVaR constraints. The manager could proceed to act as thought he or she was risk neutral constrained within those bounds.

It is difficult to make the GVaR choice as it stands outcome fully consistent with the classic Von- Neumann ~ Morgenstern choice theory under risk. The problem arises because the effective or implied GVaR utility function is not monotonically increasing. It is open to the user to construct an explicit utility function in way that captures essential features of the GVaR constraints in the form of user-assigned numerical penalties for their violation. However, it is possible to modify the original

CVaR constraint to make it compatible with classical utility theory. In that case the programming problem itself will allocate the penalty values.

Finally the imposition of GVaR constraints, or the equivalent utility function, raises issues about the optimal choice of portfolio hedges or other enhancements. Standard statistical hedging theory is based on least squares regression. However, a CVaR constraint implies the advisability of considering also the ordered mean difference of the proposed hedge against the base. The latter is a cumulative sample construct that harmonises with the censored mean that forms the statistical basis of the conditional value at risk. The shape and position of the hedge OMD against the portfolio to be hedged will indicate the value of the hedge, or which of several alternative proposed hedges are likely to be more effective.

References

Andersson, F., H. Mausser, D. Rosen and S. Uryasev (1999) 'Credit risk optimization with conditional value at risk', *Mathematical Programming, series B.*

Best, P. (1999) *Implementing value at risk*, Wiley.

Bowden, R.J. (2000) 'The ordered mean difference as a portfolio performance measure', *Journal of Empirical Finance*, 7,195-223.

Bowden, R.J. (2003) 'The zero capital approach to portfolio enhancement and overlay management', *Quantitative Finance*, 3, 251-61.

Bowden, R.J. (2005) 'Ordered mean difference benchmarking, utility generators and capital market equilibrium', *Journal of Business*, forthcoming.

Butler, C. (1998) *Mastering value at risk: A step by step guide to understanding and applying VaR,* Financial Times/Prentice- Hall.

Duffie, D. and J. Pan (1997) 'An overview of value at risk', *Journal of Derivatives*, 4,4-49.

Johnson, N.L. and S. Kotz (1970) *Continuous univariate distributions – 2,* New York: Houghton Mifflin.

Jorion, P. (2002) *Value at risk: The new benchmark for controlling market risk*, New York: McGraw-Hill.

JP Morgan (1994) *Introduction to Risk Metrics*, Manual, successive editions.

Krokhmal, P., J. Palmquist and S. Uryasev (2001) 'Portfolio optimization with conditional value at risk: objective and constraints', University of Florida, Dept of Industrial and Systems Engineering.

Lighthill, M.J.(1959) *An introduction to Fourier Analysis and generalised functions*, New York: Academic Press.

Mann, N. R., R.E. Schafer and N. Singpurwalla (1974) *Methods for the statistical analysis of reliability and life data*, New York: Wiley.

Rockafellar, R.T. (1968) 'Duality in nonlinear programming', in *Mathematics of the decision sciences*, part I, (ed.) G.B. Dantzig and A.F. Veinott, Jr., Providence R.I.:American Mathematical Society.

Rockafellar, R.T.and S. Uryasev (2000) 'Optimization of conditional value at risk', *Journal of Risk*, 2, 21-41.

Appendix I

Illustrating the 'dual' optimisation problem

(Section 2.3 refers)

Reproducing problem (2b):

$$\min_{z,r_L} r_L$$

subject to

$$(i) \quad F(r_L) = \alpha$$

$$(ii) \quad E[R] \geq m$$

$$(iii) \quad E[R \mid R \leq r_L] \geq v$$

Figure 10 illustrates with three portfolios A, B. Consider the constraints:

Portfolio A satisfies (i)-(iii);

Portfolio B satisfies (i)&(ii) but not (iii);

Portfolio C satisfies (i)-(iii).

So the choice reduces to A and C, of which A is preferred because it has the smaller value at risk r_L .

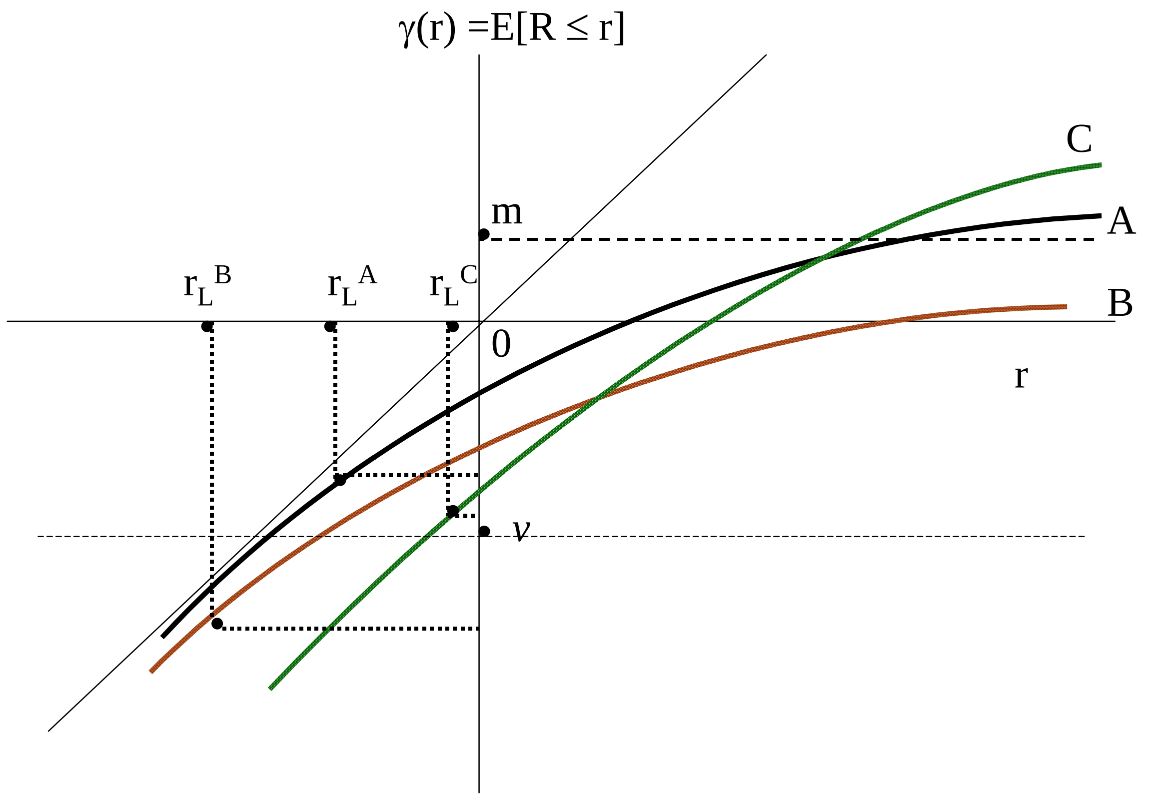

Figure 10: Minimising the value at risk

Appendix II

(Section III refers).

The logistic distribution admissible set

Proposition:

If R has a logistic distribution, the censored mean function at R = r is given by

$$\gamma(r) = r + \frac{\beta}{F(r)}\log(1 - F(r)].$$

Proof:

It is convenient to first establish the same result for the standardised distribution. Using the logistic property $f_S(Y) = F_S(Y)(1 - F_S(Y))$,

$$\Phi_S(y) = \int_{-\infty}^{y} F_S(Y)dY = \int_{-\infty}^{y} \frac{f_S(Y)}{1-F_S(Y)}dY = -\log(1-F_S(y)).$$

Hence from equation (5) of the text,

$$\gamma_S(y) = y - \frac{\Phi_S(y)}{F_S(y)} = y + \frac{1}{F_S(y)}\log(1-F_S(y)).$$

Reverting to the original scale via $y = \dfrac{r-\mu_R}{\beta}$ it follows from Lemma 1 that

$$\gamma(r) = \mu_R + \beta\gamma_S(y)$$

and the desired result is a consequence.

□

To explore the admissible set, it is convenient to work in terms of

$$x = F_S(y_L);\; y = \frac{r_L-\mu_R}{\beta};\; \beta = \sigma_R\sqrt{3}/\pi.$$

Recasting Proposition 2, the admissible set can be defined as follows:

CVaR:

(A22a) $e^{-\frac{(r_L-v)}{\beta}x} \le 1-x$

VaR:

(A22b) $x \le \alpha$.

For the CVaR condition to hold at all we must have some point of intersection between the two functions in (A22a).

Figure 11 sketches the various possibilities. Fix the significance level α and consider the effect of varying the variance-related parameter β. For the given value of α there is some value β_α and associated STD σ_α such that equality holds in (A22a). For this value the VaR and CVaR constraint hold simultaneously, i.e. both are binding. For $\beta < \beta_\alpha$, the VaR constraint will binding, but not the CVaR. For $\beta > \beta_\alpha$, things are reversed - the CVaR constraint is binding but the VaR is not. For $\beta > \beta_{max}$, there is no solution that satisfies both constraints.

The critical value is given by $\beta_{\max} = r_L - v$. This gives us an upper bound for the allowable portfolio standard deviation as

$$\sigma_R \le \frac{\pi(r_L-v)}{\sqrt{3}} \approx 1.814(r_L-v),$$

which is expression (10) of the text.

The VaR constraint (A22b) translates to the (μ_R ,σ_R) plane as the linear constraint:

$$\mu \geq r + \frac{\sqrt{3}}{\pi}\log(\frac{1-\alpha}{\alpha})\sigma_R .$$

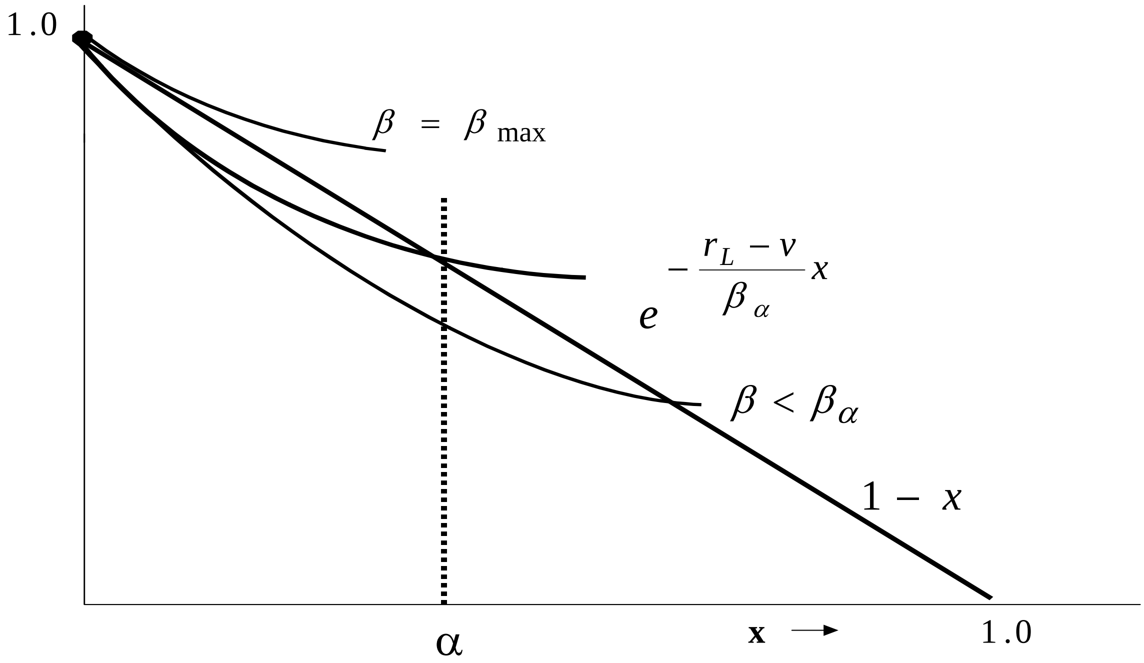

Figure 11: Existence of the logistic admissible set

Appendix III

(Section V refers).

Proof of Proposition 2 of the text.

Define the reverse hazard function $h_b(r) = \frac{f_b(r)}{F_b(r)}$ *where* $F_b(r)$ *is the distribution function of* R_b *evaluated at* $R_b = r$, *and* $f_b(r)$ *is the corresponding density function evaluated at the same point. Then*

(21)

$$\tau_{a/b} = \frac{E[R_a - R_b] + F_b(r)\{\mu E[(R_a - R_b)|R_b \le r_L] + \pi E[(R_a - r_L)|R_b = r_L]h(r_L)\}}{1 + F_b(r_L)[\mu + \pi h(r_L)]}$$

Proof:

The shortest demonstration is by using the generalised calculus associated with step and Dirac delta functions (see e.g. Lighthill 1959). We have for any variable *R*,

$$\frac{dSF(r_L - R)}{dR} = -\delta(R - r_L)\,,$$

where $\delta(x)$ is the Dirac delta function. Thus

(A23) $U'(R_b) = 1 + \mu SF(r_L - R_b) + [\pi - \mu(R_b - r_L)]\delta(R - r_L)$.

For any measurable smooth function *g(x)* (see Lighthill 1959), we have $\int_{-\infty}^{\infty} \delta(x - c)g(x)dx = g(c)$, the 'filtering property'. It follows that

$$E[U'(R_b)] = 1 + \mu F_b(r_L) + \pi f_b(r_L)\,,$$

which gives the denominator of (21).

From expression (A23),

(A24)

$$E[R_a - R_b)U'(R_b)] = E[R_a - R_b] + \mu E[(R_a - R_b)SF(r_L - R_b)] + \pi E[(R_a - R_b)\delta(R_b - r_L)]$$

Now

(A25) $E[(R_a - R_b)SF(R_b - r_L)] = E[(R_a - R_b)|R_b \le r_L]F_b(r_L)$.

Similarly, writing the joint density $f(R_a, R_b) = f(R_a|R_b) f_b(R_b)$ and using the filtering property of the Dirac delta, we get

(A26) $E[(R_a - R_b)\delta(R_b - r_L)] = E[(R_a - r_L)|R_b = r_L] f_b(r_L)$.

Substituting (A25), (A26) into (20) of the text gives the numerator of the required result (21).

□