Time Consistent Dynamic Risk Measures

Kang Boda, Jerzy A. Filar

Center for the Industrial and Applied Mathematics,

School of Mathematics & Statistics,

University of South Australia,

Mawson Lakes, SA 5095, Australia

Outline

1. Measures of Risk
2. Time Consistency
3. Target Percentile Risk Measures
4. Multi-Stage VaR and CVaR

Measure of Risk

- Variance (or, equivalently, standard deviation) – Markowitz's seminar work.
- "Value-at-risk" (VaR) – the maximum loss that might be incurred with respect to a given, and fixed, confidence level.
- Rockafellar and Uryasev, "Conditional value-at-risk" (CVaR) – Expected losses exceeding VaR shown in the following picture:

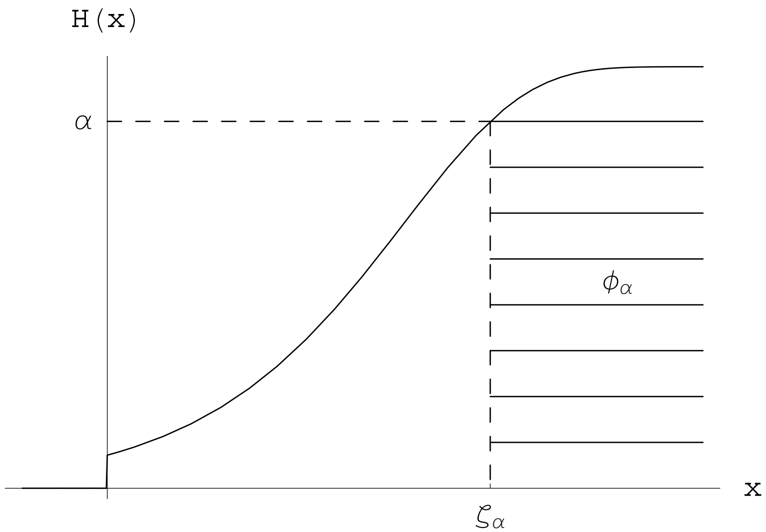

Why multi-stage?

- Static (single-stage), portfolio allocation.
- Nowadays, most investors are making portfolio decisions dynamically (over time); usually at discrete times (e.g., once a day or once a week).
- Inspired by the "principle of optimality" of dynamic programming, we believe that "*time consistency*" is an important property of multi-stage measure of risk.
- Multi-stage VaR and CVaR need not be time consistent.
- *What measure of risk is most appropriate for a dynamic portfolio allocation problem?*

Notations

An $n-$stage portfolio optimization planning problem is considered.

- Decision rule $\pi_t, t = 1, \ldots, n$ — a column vector, whose entries represent the fractions of the total portfolio allocated to individual stocks.
- A *policy* $\pi = (\pi_1, \cdots, \pi_n)$ — a sequence of decision rules.
- Random variable r_t — the return/reward at stage t.
- $R_t := g(r_1, \cdots, r_t)$, — the aggregated return.
- $Z_t^\pi = Z_t(\pi_1, \cdots, \pi_t)$ — the risk measure of $t-$stage problem – with respect to the function g – corresponding to the decision rules $\pi_1, \cdots, \pi_t$.

Time consistent risk measures

The risk measure Z is *time-consistent* if two conditions are satisfied:

TC1 If the decision rule π_t^* at each stage t, $t = 1, \cdots, n$ is chosen by

$$\pi_t^* \in \text{Argmin}_{\pi_t} Z_{n-t+1}(\pi_t, \pi_{t+1}^*, \cdots, \pi_n^*), \ \forall t = 1, \cdots, n;$$

then the policy $\pi^* = (\pi_1^*, \cdots, \pi_n^*)$ will be the optimal policy in the problem

$$\min_{\pi} Z_n(\pi).$$

Motivation: **TC1** ensures that a policy assembled from risk measure minimizing decision rules, in"reverse time", is a risk measure minimizing policy over the entire horizon.

Time consistent risk measures Cont.

TC2 If the policy is optimal over the entire horizon, that is,

$$\pi^* = (\pi_1^*, \cdots, \pi_n^*) \in \mathrm{Argmin}_\pi \, Z_n(\pi),$$

then it also satisfies

$$(\pi_t^*, \cdots, \pi_n^*) \in \mathrm{Argmin}_{\pi_t, \cdots, \pi_n} Z_{n-t+1}(\pi_t, \cdots, \pi_n), \ \forall t = 2, \cdots, n.$$

Motivation: **TC2** ensures that “sub-policies” of an optimal policy π^*, over the remaining (shorter than n) time horizons, are also risk minimizing policies in the corresponding (shorter) sub-problems.

VaR and CVaR need not be time consistent

- In our paper,

  Kang Boda and Jerzy A.Filar, Time Consistent Dynamic Risk Measures, to appear in MMOR **62(3)**, 2005,

  http://www.math.vu.nl/ koole/MMOR/

  we give an example showing that multi-stage VaR and CVaR are **NOT** time consistent.

- This raises a natural question:

  Can we construct a time consistent risk measure?

Target percentile risk measures

Discrete-time and stationary *Markov decision process* (MDP):

$$\Gamma = (S, A, R, P, \beta)$$

where

- S — *State space*, countable.
- $A(i), i \in S$ — *action space* in each state i, finite.
- $A = \cup_{i \in S} A(i)$ — *overall action space*, countable.
- R — *return set*, bounded and countable.
- $p_{ijr}^a := P(i_{t+1} = j, r_t = r | i_t = i, a_t = a),\ i, j \in S, a \in A(i), r \in R, n \geq 1$ — single stage conditional *transition probabilities*.
- $\beta \in (0, 1)$ — *discount factor*.

The extended (hybrid) MDP $\tilde{\Gamma}$

$$\tilde{\Gamma} = (E, A, R, P, \beta)$$

where

- $E = S \times \mathcal{R}$ — state space, e.g. $e = (i, x)$=(old state, target).
- $A = \cup_{(i,x)\in E} A(i, x) = \cup_{i \in S} A(i)$ — action space and $A(i, x) = A(i), (i, x) \in E$.
- P — the extended conditional transition probabilities:

$$P(e_{t+1} = (j, \tfrac{x-r}{\beta}) | e_t = (i, x), a_t = a) = p^a_{ijr},$$
$$i, j \in S, a \in A(i), r \in R, x \in \mathcal{R}.$$

- The return set R and the discount factor β are the same as in MDP Γ.

Policies (Controls)

The *admissible history* up to stage t,

$$h_t = (i_1, x_1, a_1, \cdots, i_{t-1}, x_{t-1}, a_{t-1}, i_t, x_t)$$

where $(i_k, x_k) \in E$, $a_k \in A(i_k, x_k)$, $k = 1, \cdots, t-1$, $(i_t, x_t) \in E$.
Based on the above histories, we have:

- Π_m — set of all *Markov policies*.
- Π_m^d — set of all *deterministic Markov policies*.
- Π_s — set of all *stationary policies*.
- Π_s^d — set of all *deterministic stationary policies*.
- Π_0 — set of all *TI-policies*, the policy $\pi \in \Pi_0$ is independent of all targets x_t $(t \geq 1)$.

Percentile Optimality Criterion

Sum of discounted returns generated by policy π for the n-stage finite horizon problem:

$$R_n^\pi = \sum_{t=1}^{n} \beta^{t-1} r_t, \; n \geq 1.$$

Our objective function, for any $\pi \in \Pi$ will be:

$$F_n^\pi(i,x) = P_\pi(R_n^\pi \leq x | e_1 = (i,x)), \; (i,x) \in E, n \geq 1,$$

that the decision maker wishes to minimize if he or she is interested in letting the return as big as possible.

Our optimal value function will be:

$$F_n^*(i,x) = \inf_{\pi \in \Pi} \{F_n^\pi(i,x)\}, \; (i,x) \in E, n \geq 1.$$

Percentile Optimality Criterion Cont.

$\pi^* \in \Pi$ is called an *n-stage optimal policy* if:

$$F_n^{\pi^*}(i,x) = F_n^*(i,x), \forall\,(i,x) \in E, n \geq 1.$$

We have the following theorem on the existence of optimal policies.

Theorem 0.1 *(i) The optimal value function* $\{F_n^*, n \geq 0\}$ *satisfies the following optimality equations,*

$$F_0^* = I_{[0,\infty)},\; F_n^* = TF_{n-1}^*, \quad n \geq 1;$$

where T is a suitably defined "dynamic programming" operator.

(ii) For all $n \geq 0$, $i \in S$, $F_n^*(i,x)$ is a distribution function of some random variable X taking on values x;

(iii) For all $n \geq 0$, there exists a policy $\pi \in \Pi_m^d$ such that $F_n^\pi = F_n^*$.

Time consistency of F_n^π

Optimal action sets are defined by:

$$A_n^*(i,x) := \{a | a \in A(i) \text{ and } \ KF_{n-1}^*(i,x,a) = F_n^*(i,x)\},$$
$$(i,x) \in E, n \geq 1,$$

where K is a standard one step minimization operator.

It follows from Theorem 0.1 that $A_n^*(i,x) \neq \emptyset$, $\forall e = (i,x) \in E,\ n \geq 1$.

Lemma 0.1 $\forall e \in E,\ k = 1, \cdots, n,\ \delta_k(e) \in A_k^*(e) \Longrightarrow$ $\pi(n) = (\delta_n, \delta_{n-1}, \cdots, \delta_1)$ *is n-stages optimal.*

Lemma 0.1 concludes that a policy assembled from risk measure minimizing decision rules, in "reverse time", is a risk measure minimizing policy over the entire horizon, which shows that F_n^π satisfies **TC1**.

Time consistency of F_n^π Cont.

Definition: $\bar{\pi}^{(i,x,a)} = (\pi_t^{(i,x,a)}, t \geq 1)$ is called a *"cut-head policy"* if

$$\pi_t^{(i,x,a)}(\cdot|h_t) = \pi_{t+1}(\cdot|i,x,a,h_t), \forall\, h_t \in H_t,\ t \geq 1.$$

Theorem 0.2 *Let $\pi = (\pi_k, k \geq 1) \in \Pi$, for a given $(i,x) \in E$, then $F_n^\pi(i,x) = F_n^*(i,x)$ if and only if $\pi_1(A_n^*(i,x)|i,x) = 1$ and*

$$F_{n-1}^{\bar{\pi}^{(i,x,a)}}(j,(x-r)/\beta) = F_{n-1}^*(j,(x-r)/\beta),$$

whenever $\pi_1(a|i,x)p_{ijr}^a > 0$.

Theorem 0.2 tells us that the "sub-policies" of an optimal policy π^*, over the remaining (shorter than n) time horizons, are also risk minimizing policies in the corresponding (shorter) sub-problems.

*Conclusion: The objective function $F_n^\pi(i,x)$ in the target percentile MDP satisfies time consistency condition (**TC1** and **TC2**).*

Complete Stochastic Order Optimization

The preceding time-consistent measure of risk $F_n^\pi(i, x)$ depends on the target level x. Can we eliminate this dependence?

Now we consider the following discrete time and stationary MDPs,

$$\Gamma^0 = (S, A, p, R)$$

where

- S, $A = \bigcup_{i \in S} A(i)$ — the state space and action space, both are countable.
- $A(i), i \in S$ — the set of admissible actions when the system is in state i, finite.
- $p = (p_{ij}^a;\ i, j \in S, a \in A(i))$, $r = r(i, a, j), i, j \in S, a \in A(i)$ — the conditional transition probabilities and the return function.

New Optimality

Sum of discounted returns generated by policy $\pi \in \Pi_0$ that won't depend on the target x for the n-stage finite horizon problem:

$$R_n^\pi = \sum_{t=1}^{n} \beta^{t-1} r_t,\ n \geq 1.$$

Now our *objective function*, for any $\pi \in \Pi_0$ will be:

$$V_n^\pi(i,x) := P_\pi(R_n^\pi \leq x | i_1 = i), \forall i \in S, x \in \mathcal{R}, \pi \in \Pi_0, n \geq 1;$$

that the decision maker wishes to minimize if he or she is interested in letting the return as big as possible, conditioned only on $i_1 = i$, not on $e_1 = (i,x)$.

Our new *optimal value function* will be:

$$V_n^*(i,x) := \inf_{\pi \in \Pi_0} \{V_n^\pi(i,x)\}, i \in S, x \in \mathcal{R}, n \geq 1.$$

New Optimality Cont.

$\pi^* \in \Pi_0$ is called an *n-stage optimal policy* if:

$$V_n^{\pi^*}(i,x) = V_n^*(i,x),\ \forall i \in S, x \in \mathcal{R}.$$

Theorem 0.3 *For a given target x and any $n \geq 1$, there exists a $\pi^* \in \Pi_0$ such that,*

$$V_n^{\pi^*}(i,x) = V_n^*(i,x), \quad \forall i \in S.$$

Unfortunately, the new objective function $V_n^{\pi}(i,x)$ is no longer time consistent, except under additional strong conditions.

Multi-stage VaR

Definition

For a given policy $\pi \in \Pi_0$, initial state $i \in S$ and the probability threshold $\alpha \in [0,1]$, the value-at-risk $(\zeta^{\pi}_{n,\alpha}(i))$ for the n-stage return R^{π}_n is denoted by:

$$\zeta^{\pi}_{n,\alpha}(i) := -\sup\{\zeta | V^{\pi}_n(i,\zeta) \leq \alpha\}, \forall i \in S, \pi \in \Pi_0, \alpha \in [0,1], n \geq 1.$$

Motivation

Starting from an initial state i and continuing to use the policy π for $n-$stages, the decision-maker wants to minimize the loss associated with the **best** of the $100\alpha\%$ **worst** cases of the $n-$stage total return R^{π}_n, namely, $\zeta^{\pi}_{n,\alpha}(i)$.

Optimality

Optimal value functions

$$\zeta^*_{n,\alpha}(i) := \inf_{\pi \in \Pi_0} \{\zeta^\pi_{n,\alpha}(i)\}, \forall i \in S, \alpha \in [0,1], n \geq 1. \tag{1}$$

1. $\alpha-$ *optimal policy* $\pi^\alpha \in \Pi_0$ for $n-$ stage value-at-risk $(\zeta^\pi_{n,\alpha}(i))$ is
$$\zeta^{\pi^\alpha}_{n,\alpha}(i) = \zeta^*_{n,\alpha}(i), \forall i \in S, n \geq 1.$$
2. *Optimal policy* $\pi^* \in \Pi_0$ for $n-$stage value-at-risk $(\zeta^\pi_{n,\alpha}(i))$ is
$$\zeta^{\pi^*}_{n,\alpha}(i) = \zeta^*_{n,\alpha}(i), \forall i \in S, \alpha \in [0,1], n \geq 1.$$

We shall show that, with the help of our target percentile formulation, the optimization problem implied by (1) is tractable.

Existence of Multi-stage VaR

Define $x^*_{n,\alpha}(i)$ with respect to the optimal value function:

$$x^*_{n,\alpha}(i) := -\sup\{x | V^*_n(i,x) \leq \alpha\}, i \in S, \alpha \in [0,1], n \geq 1.$$

Theorem 0.4 *For all* $i \in S, \alpha \in [0,1]$, *and* $n \geq 1$, *we have*

$$x^*_{n,\alpha}(i) = \zeta^*_{n,\alpha}(i) = \inf_\pi \zeta^\pi_{n,\alpha}(i),$$

and there exists an $\alpha-$*optimal policy* $\hat{\pi}_\alpha \in \Pi_0$ *such that*

$$\zeta^{\hat{\pi}_\alpha}_{n,\alpha}(i) = \zeta^*_{n,\alpha}(i).$$

If $A^*_k(i) = \cap_{x \in \mathcal{R}} A^*_k(i,x) \neq \emptyset, \forall i \in S, k = 1, \cdots, n$, *then there exists an* $n-$*stage optimal, time consistent policy* $\hat{\pi} \in \Pi_0$ *such that*

$$\zeta^{\hat{\pi}}_{k,\alpha}(i) = \zeta^*_{k,\alpha}(i), \ \forall k, \alpha.$$