Numerical Solution of the Hamilton-Jacobi-Bellman Formulation for Continuous Time Mean Variance Asset Allocation

Peter Forsyth[1] Jian Wang[2]

[1]Cheriton School of Computer Science
University of Waterloo

[2]Royal Bank of Canada, Toronto

Minneapolis, October 7, 2013

Objective of this talk

- We develop numerical methods for solving HJB PDEs for
  - Pre-commitment mean-variance
  - Time consistent mean-variance
  - Mean-quadratic variation
- We can apply arbitrary constraints to the optimal policy
  - Example: investor who seeks to maximize wealth/income ratio at retirement
- With realistic constraints (no bankruptcy, maximum leverage ratio)
  - All efficient frontiers are very similar
  - But optimal policies are very different

Defined Contribution Pension Plan Wealth Accumulation

We consider an investor who seeks to maximize the ratio of accumulated wealth to their annual salary at retirement (Cairns et al, 2006).

Investor can allocate wealth to risky asset S which follows

$$
\begin{aligned}
dS &= (r + \xi\sigma)S\ dt + \sigma S\ dZ_1 \\
& \quad r = \text{ risk free rate } \ ; \ \ dZ_1 = \text{ increment of a Wiener process} \\
& \quad \sigma = \text{ volatility } \ ; \ \ \xi = \text{ market price of risk}
\end{aligned}
$$

Or to a risk-free asset B which follows

$$
\frac{dB}{dt} = rB
$$

The investor's yearly stochastic salary Y follows

$$
\begin{aligned}
dY &= (r + \mu_Y)Y\ dt + \sigma_{Y_0} Y\ dZ_0 + \sigma_{Y_1} Y\ dZ_1 \\
& \quad dZ_0 \text{ is independent of } \ dZ_1
\end{aligned}
$$

Wealth to Income Ratio: II

The investor contributes into the pension plan account at a rate of πY per year.

The investor allocates a fraction p to the risky asset S and $(1-p)$ to the riskless asset B.
The total wealth $W = S + B$ then follows

$$dW = (r + p\xi\sigma)W\,dt + p\sigma W dZ_1 + \pi Y dt$$

Define new state variable $X(t) = W(t)/Y(t)$

Then by Ito's Lemma, we obtain

$$\begin{aligned} dX &= [\pi + X(-\mu_Y + p\sigma(\xi - \sigma_{Y_1}) + \sigma_{Y_0}^2 + \sigma_{Y_1}^2)]dt \\ &\quad -\sigma_{Y_0} X dZ_0 + X(p\sigma - \sigma_{Y_1})dZ_1 \end{aligned}$$

Optimal Strategy

Define

$$\begin{aligned} p(X,t) &= \text{dynamic fraction invested in the risky asset} \\ X_T &= \text{terminal wealth to income ratio} \end{aligned}$$

Let

$$\begin{aligned} E_{t,x}^{p(\cdot)}[\cdot] &= E[\cdot|X(t)=x] \text{ with } \ p(s,X(s)), s \geq t \\ &\quad \text{being the strategy along path } \ X(s), s \geq t \\ \text{Var}_{t,x}^{p(\cdot)}[\cdot] &= \text{Var}[\cdot|X(t)=x] \text{ Variance under strategy } p(\cdot) \\ &\quad \text{along path } \ X(s), s \geq t \end{aligned}$$

Mean Variance Optimization

Now, we seek to find the control $p(\cdot)$ which solves [1]

$$\sup_{p(\cdot)\in\mathbb{P}} \left\{ \underbrace{E_{t,x}^{p(\cdot)}[X_T]}_{Reward} - \lambda \underbrace{Var_{t,x}^{p(\cdot)}[X_T]}_{Risk} \right\} ,$$
$$\mathbb{P} \text{ is the set of admissible controls} \tag{1}$$

Solving the above problem for $\lambda \in [0, \infty)$ traces out the efficient frontier.

[1] Using a scalarization method to determine the Pareto optimal points

LQ Embedding (Zhou and Li (2000), Li and Ng (2000))

Equivalent formulation: for fixed λ, if $p^*(\cdot)$ maximizes

$$\sup_{p(\cdot)\in\mathbb{P}} \Big\{ \underbrace{E_{t,x}^{p(\cdot)}[X_T]}_{Reward} - \lambda \underbrace{Var_{t,x}^{p(\cdot)}[X_T]}_{Risk} \Big\} \, ,$$

$$\mathbb{P} \text{ is the set of admissible controls} \tag{2}$$

then there exists a $\gamma = \gamma(t, x, E[X_T])$ such that $p^*(\cdot)$ minimizes [2]

$$\inf_{p(\cdot)\in\mathbb{P}} E_{t,x}^{p(\cdot)}\left[\left(X_T - \frac{\gamma}{2} \right)^2 \right] . \tag{3}$$

$\hookrightarrow$ Equation (3) can be solved using dynamic programming.

---

[2] Strictly speaking, since some values of γ may not represent points on the original frontier, we need to use the algorithm in Tse, Forsyth, Li (2012) to remove these points.

Wealth to Income Ratio: III

Let $V(x,\tau) = E^{p(\cdot)}_{\tau=T-t,x}[(X_T - \gamma/2)^2]$, then usual steps determine the HJB equation for $V(x,\tau)$

$$
\begin{aligned}
V_\tau &= \inf_{p\in\mathbb{P}}\left\{ \mu_x^p V_x + \frac{1}{2}(\sigma_x^p)^2 V_{xx} \right\} \\
& \mu_x^p = \pi + x(-\mu_Y + p\sigma(\xi_1 - \sigma_{Y_1}) + \sigma_{Y_0}^2 + \sigma_{Y_1}^2) \\
& (\sigma_x^p)^2 = x^2(\sigma_{Y_0}^2 + (p\sigma - \sigma_{Y_1})^2) \\
& V(x,\tau=0) = (x - \frac{\gamma}{2})^2
\end{aligned}
$$

We can easily develop numerical methods for solving this HJB equation, for arbitrary constraints on the control.

- We can guarantee convergence to the viscosity solution

Pre-Commitment Policy: Not Time Consistent

But

$$
\begin{aligned}
p_t^*(s, X(s)) \quad &\neq \quad p_{t+\Delta t}^*(s, X(s)) \; ; \; s \geq t + \Delta t \; , \\
& \quad p_t^* = \text{ Optimal Policy as seen at time } t \\
& \quad p_{t+\Delta t}^* = \text{ Optimal Policy as seen at time } t + \Delta t
\end{aligned}
$$

The strategy which solves problem (2) has been called the *pre-commitment* policy (Basak,Chabakauri: 2010; Bjork et al: 2010)

- Much discussion on the economic meaning of such strategies.
- However, there are economic situations where pre-commitment is the correct strategy
  - Optimal trade execution
  - Asset liability management (if I have time, I will relate a personal story about this)
- Different applications may require different strategies.

If you don't like pre-commitment: How about Mean Quadratic Variation?

Formally, the quadratic variation risk measure is defined as

$$\text{Risk}_t = E\left[\int_t^T \left(dX(t')\right)^2 \right]$$

$$X = \text{ wealth to income ratio}$$

This is the quadratic variation of the $X(t) = W(t)/Y(t)$ process.

This measures risk in terms of the average variability of wealth (in units of yearly income) along the entire trading path.

This problem is naturally time consistent.

Mean Quadratic Variation

Find optimal strategy $p(\cdot)$ which maximizes (for fixed λ)

$$\sup_{p(\cdot)\in\mathbb{P}} \left\{ \underbrace{E_{x,t}^{p(\cdot)}\left[B_T\right]}_{Reward} - \lambda \underbrace{E_{x,t}^{p(\cdot)}\left[\int_t^T \left(dX(t')\right)^2\right]}_{Risk} \right\} \tag{4}$$

Originally suggested as a risk measure by Brugierre (1996).

Industry standard approach to optimal trade execution uses quadratic variation as the risk measure.[3]

One can easily derive the HJB equation for the optimal control $p^*(\cdot)$ for Mean Quadratic variation optimal strategies

Varying λ will trace out a curve in the expected value, quadratic variation plane

[3] Although this is not widely known. They think it's a good approximation to variance (Almgren and Chriss (2001))

Time consistent mean-variance

Another alternative to pre-commitment: maximize (for fixed λ)

$$\sup_{p(\cdot)\in\mathbb{P}} \left\{ \underbrace{E^{p}_{t,x}[X_T]}_{Reward} - \lambda \underbrace{Var^{v}_{t,x}[X_T]}_{Risk} \right\} ,$$

$$\mathbb{P} \text{ is the set of admissible controls} \tag{5}$$

But we require that controls are time consistent

$$p^*_t(s, X(s)) = p^*_{t+\Delta t}(s, X(s)) \; ; \; s \geq t + \Delta t \; , \tag{6}$$

Suggested by (Basak,Chabakauri: 2010; Bjork et al: 2010).

Can devise numerical method for solving this problem (EJOR: Wang, Forsyth (2011))

Time Consistent Mean Variance

$$\sup_{p(\cdot)\in\mathbb{P}} \Big\{ \underbrace{E_{t,x}^{p(\cdot)}[X_T]}_{Reward} - \lambda \underbrace{Var_{t,x}^{p(\cdot)}[X_T]}_{Risk} \Big\} \ ,$$
$$p_t^*(s, X(s)) = p_{t+\Delta t}^*(s, X(s)) \ ; \ s \geq t + \Delta t \qquad (7)$$

But:

- Why assume $\lambda = const.$?
- Is this Pareto optimal in any sense?
- if $\lambda = \lambda(X) \Rightarrow$ strange results
- In the unconstrained control case
  - $\rightarrow$ For the case of nonstochastic salary, the optimal policy for time consistent mean variance is identical to the optimal policy for mean quadratic variation (closed form solution available)

Admissible Controls

Recall that $\mathbb{P}$ is the set of admissible controls.
Let $\mathbb{D}$ be the set of possible values of x (wealth/income)
We consider the following cases:

Case	$\mathbb{D}$	$\mathbb{P}$
Allow Bankruptcy	$[-\infty, +\infty]$	$(-\infty, +\infty)$
No Bankruptcy	$[0, \infty]$	$[0, +\infty)$
Bounded Control	$[0, \infty]$	$[0, p_{\max}]$

- Bankruptcy: trading continues even if insolvent (not realistic)
- No bankruptcy: also not shorting stock
- Bounded control: realistic case, i.e. $p_{\max} = 1.5$ corresponds to maximum borrowing of 50% of net wealth

Convergence to the Viscosity Solution: Mean Variance Pre-commitment case; Finite Difference Method

Theorem

Provided that (a) the HJB PDE satisfies a strong comparison result (b) fully implicit timestepping is used with a positive coefficient discretization, then the discrete scheme converges to the unique, continuous viscosity solution of the control problem.

Proof.

The scheme is

- Unconditionally monotone
- Unconditionally stable
- Consistent in the viscosity sense

Hence convergence follows from the results in (Barles, Souganidis (1991)). □

Wealth to Income Ratio Case

μ_y	0.	ξ_1	0.2
σ	0.2	σ_{Y1}	0.05
σ_{Y0}	0.05	π	0.1
T	20 years		
$\mathbb{P}$	$[0, 1.5]$	$\mathbb{D}$	$[0, +\infty)$

The total wealth W follows

$$dW = (r + p\xi\sigma)W\ dt + p\sigma W dZ_1 + \pi Y dt$$

The investor's yearly salary Y follows

$$dY = (r + \mu_Y)Y\ dt + \sigma_{Y_0} Y\ dZ_0 + \sigma_{Y_1} Y\ dZ_1$$
$$dZ_0 \text{ is independent of } dZ_1$$

We seek to maximize $X(T) = W(T)/Y(T)$.
Risky asset proportion $p \in \mathbb{P}$, $x = w/y \in \mathbb{D}$.

Efficient Frontiers: Pre-commitment mean variance

- No analytic solutions available
- No risk free portfolio (due to salary risk)
- $X(t = 0) = .5$

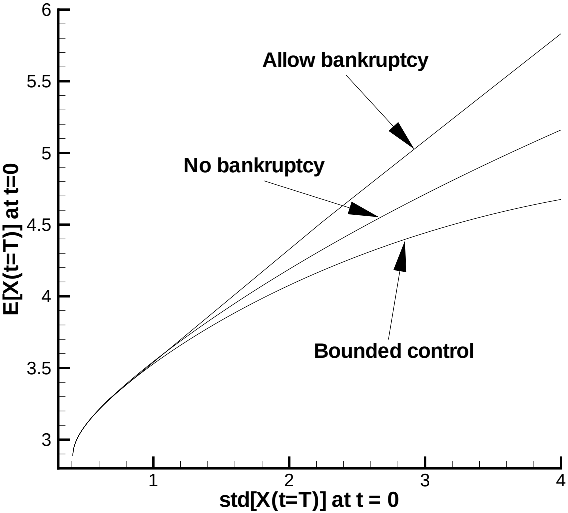

Optimal Strategy: Pre-commitment mean variance

- $X(t=0) = 0.5$, $(\text{Std}^{p^*}_{t=0}[X_T], E^{p^*}_{t=0}[X_T]) = (1.7407, 3.9551)$.

- As time goes on, and if x has not increased, the investor takes on more risk (agrees with analytic solution, non-stochastic salary case)
- At any time, if (W/Y) is large enough, all wealth is switched to the risk free asset

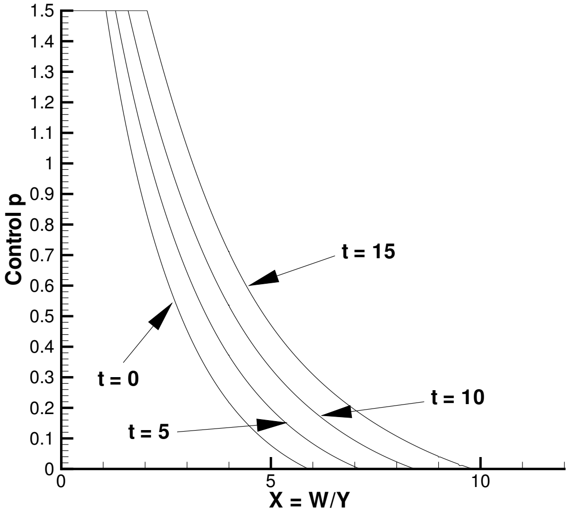

Efficient Frontiers, Bounded Control: comparison of strategies

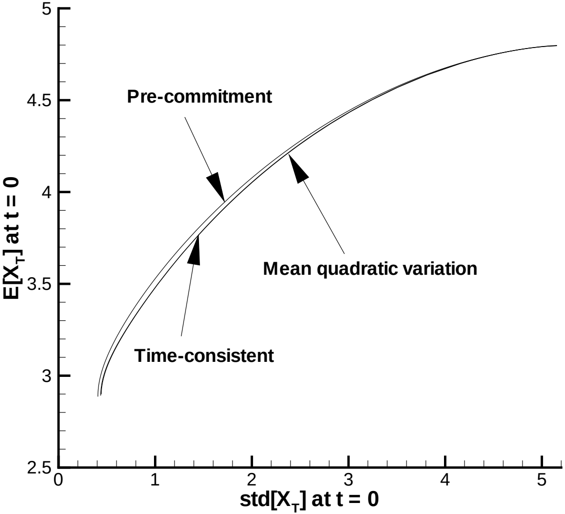

Comparison of Controls

- $(X(t = 0) = 0.5$
- Control at $t = 0+$
- $\text{Std}_{t=0,x}^{p^*}[X(T)] \simeq 3.24$
- Efficient Frontiers similar
- Controls quite different

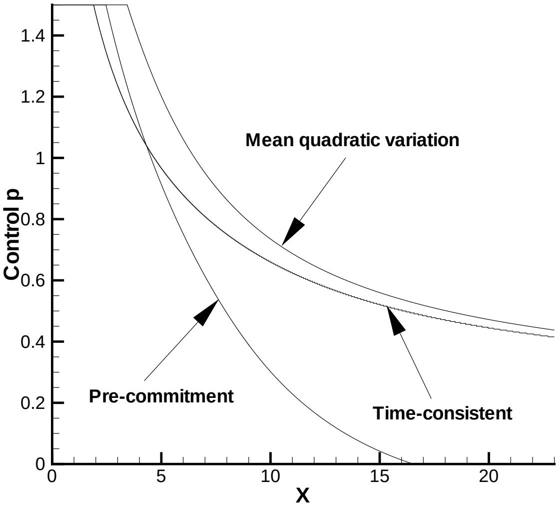

Conclusions

- There exists an HJB equation which generates the optimal control for continuous time mean-variance portfolio optimization (*pre-commitment policy*)
- This HJB equation can be reliably solved, and used to generate points on the efficient frontier
  - Any type of constraint can be applied to the control
  - Assuming a strong comparison property holds, then the scheme will converge to the viscosity solution of the HJB PDE
- Similar techniques can be used to generate controls for
  - Time consistent mean variance
  - Mean quadratic variation
- For bounded control case, all these strategies give similar efficient frontiers, but controls are different
- Which strategy should we choose?
  - Sometimes, pre-commitment is optimal (college fund for my son)
  - Can extend numerical PDE method to *Better than mean variance* (Cui, Li, Wang, Zhu, 2012) strategy (jump diffusion)