Better than pre-commitment mean-variance portfolio allocation strategies: a semi-self-financing Hamilton-Jacobi-Bellman equation approach

Peter Forsyth[1] D.M. Dang[1]

[1]Cheriton School of Computer Science
University of Waterloo

Chinese University of Hong Kong, July 8, 2014

The Basic Problem

Many financial problems have unhedgeable risk

- Optimal trade execution: Broker sells large block of shares
  - $\rightarrow$ Maximize average price, minimize risk, taking into account price impact
- Long term asset liability management (insurance)
  - $\rightarrow$ Match liabilities with minimal risk
- Minimum variance hedging of contingent claims (with real market constraints)
  - $\rightarrow$ Liquidity effects, different rates for borrowing/lending
- Pension plan investments.
- Wealth management products

Risk-reward tradeoff

All these problems (and many others) involve a tradeoff between risk and reward.

- A classic approach is to use some sort of utility function
- But this has all sorts of practical limitations
  - → What is the utility function of an investment bank?
  - → What *risk aversion* parameter should be selected by the Pension Investment Committee?

Alternative: mean-variance optimization

- When risk is specified by variance, and reward by expected value
  - → Even non-technical managers can understand the tradeoffs and make informed decisions[1]

[1] I am now a member of the University of Waterloo Pension Committee. I have seen this problem first-hand

Multi-period Mean Variance

Some issues:

- Standard formulation not amenable to use of dynamic programming
- Criticism: variance as risk measure penalizes upside as well as downside
- Pre-commitment mean variance strategies are not time consistent

I hope to convince you that multi-period mean variance optimization is

- Intuitive
- Can be modified slightly to be (effectively) a downside risk measure

Example: Wealth Management (Target Date Fund)

Example: Target Date (Lifecyle) Fund with two assets

Risk free bond B

$$dB = rB\ dt$$
$$r = \text{ risk-free rate}$$

Amount in risky stock index S

$$dS = (\mu - \lambda\kappa)S\ dt + \sigma S\ dZ + (J-1)S\ dq$$

$$\mu = \mathbb{P} \text{ measure drift} \quad ; \quad \sigma = \text{ volatility}$$
$$dZ = \text{ increment of a Wiener process}$$

$$dq = \begin{cases} 0 & \text{with probability } 1 - \lambda dt \\ 1 & \text{with probability } \lambda dt, \end{cases}$$
$$\log J \sim \mathcal{N}(\mu_J, \sigma_J^2). \quad ; \quad \kappa = E[J-1]$$

Optimal Control

Define:

$$\begin{aligned} X &= (S(t), B(t)) = \text{Process} \\ x &= (S(t) = s, B(t) = b) = (s, b) = \text{State} \\ (s+b) &= \text{total wealth} \end{aligned}$$

Let $(s, b) = (S(t^-), B(t^-))$ be the state of the portfolio the instant before applying a control

The control $c(s, b) = (d, B^+)$ generates a new state

$$\begin{aligned} b &\rightarrow B^+ \\ s &\rightarrow S^+ \\ &S^+ = \underbrace{(s+b)}_{\text{wealth at } t^-} - B^+ - \underbrace{d}_{\text{withdrawal}} \end{aligned}$$

Note: we allow cash withdrawals of an amount d at a rebalancing time

Semi-self financing policy

Since we allow cash withdrawals

$\rightarrow$ The portfolio may not be self-financing

$\rightarrow$ The portfolio may generate a free cash flow

Let $W_a = S(t) + B(t)$ be the **allocated wealth**

- W_a is the wealth available for allocation into $(S(t), B(t))$.

The non-allocated wealth $W_n(t)$ consists of cash withdrawals and accumulated interest

Constraints on the strategy

The investor can continue trading only if solvent

$$\underbrace{W_a(s,b) = s + b > 0}_{Solvency\ \ condition}. \tag{1}$$

In the event of bankruptcy, the investor must liquidate

$$S^+ = 0 \quad ; \quad B^+ = W_a(s,b) \quad ; \quad \text{if } \underbrace{W_a(s,b) \leq 0}_{bankruptcy} \ .$$

Leverage is also constrained

$$\frac{S^+}{W^+} \leq q_{\max}$$

$$W^+ = S^+ + B^+ = \text{ Total Wealth}$$

Mean and Variance under control $c(X(t), t)$

Let:

$$\underbrace{E_{t,x}^{c(\cdot)}[W_a(T)]}_{Reward} = \text{Expectation conditional on } (x, t) \text{ under control } c(\cdot)$$

$$\underbrace{\text{Var}_{t,x}^{c(\cdot)}[W_a(T)]}_{Risk} = \text{Variance conditional on } (x, t) \text{ under control } c(\cdot)$$

Important:

- mean and variance of $W_a(T)$ are as observed at time t, initial state x.

Basic Problem: Find Pareto Optimal Strategy

We desire to find the trading strategy $c^*(\cdot)$ such that, there exists no other other strategy $c(\cdot)$ such that

$$\underbrace{E_{t,x}^{c(\cdot)}[B_T]}_{Reward\ \ under\ \ strategy\ \ c(\cdot)} \geq \underbrace{E_{t,x}^{c(^*\cdot)}[B_T]}_{Reward\ \ under\ \ strategy\ \ c^*(\cdot)}$$

$$\underbrace{\text{Var}_{t,x}^{c(\cdot)}[B_T]}_{Risk\ \ under\ \ strategy\ \ c(\cdot)} \leq \underbrace{\text{Var}_{t,x}^{c^*(\cdot)}[B_T]}_{Risk\ \ under\ \ strategy\ \ c^*(\cdot)}$$

and at least one of the inequalities is strict.

In other words

- There exists no other strategy which simultaneously has higher expected value and smaller variance
- This is a Pareto optimal strategy
- There is a family of such strategies

Mean Variance: Standard Formulation

In order to solve Pareto optimization problem, we use a standard scalarization method [2].

We construct the efficient frontier by finding the optimal control $c(\cdot)$ which solves (for fixed λ)

$$\sup_{c} \Big\{ \underbrace{E_{t,x}^{c(\cdot)}[W_a(T)]}_{Reward} - \lambda \underbrace{Var_{t,x}^{c(\cdot)}[W_a(T)]}_{Risk} \Big\} \tag{2}$$

- Varying $\lambda \in [0, \infty)$ traces out the efficient frontier
- $\lambda = 0; \rightarrow$ we seek only maximize cash received, we don't care about risk.
- $\lambda = \infty \rightarrow$ we seek only to minimize risk, we don't care about the expected reward.

[2] We may not find all the Pareto optimal points by this method, unless the achievable set in the $(E^c(W_a(T)), Var^c[W_a(T)])$ plane is convex

Mean Variance: Standard Formulation

The objective is to determine the strategy $c(\cdot)$ which maximizes

$$\sup_{c(X(u),u\geq t)} \left\{ \underbrace{E_{t,x}^{c(\cdot)}[B_L]}_{Reward\ \ as\ \ seen\ \ at\ \ t} - \lambda \underbrace{Var_{t,x}^{c(\cdot)}[B_L]}_{Risk\ \ as\ \ seen\ \ at\ \ t} \right\} , \qquad \lambda \in [0, \infty) \tag{3}$$

- Let $c_t^*(x, u), u \geq t$ be the optimal policy for (3).

Then $c_{t+\Delta t}^*(x, u), u \geq t + \Delta t$ is the optimal policy which maximizes

$$\sup_{c(X(u),u\geq t+\Delta t))} \left\{ \underbrace{E_{t+\Delta t, X(t+\Delta t)}^{c(\cdot)}[B_L]}_{Reward\ \ as\ \ seen\ \ at\ \ t+\Delta t} - \lambda \underbrace{Var_{t+\Delta t, X(t+\Delta t)}^{c(\cdot)}[B_L]}_{Risk\ \ as\ \ seen\ \ at\ \ t+\Delta t} \right\} .$$

Pre-commitment Policy

However, in general

$$\underbrace{c_t^*(X(u),u)}_{optimal\ policy\ as\ seen\ at\ t} \neq \underbrace{c_{t+\Delta t}^*(X(u),u)}_{optimal\ policy\ as\ seen\ at\ t+\Delta t} \quad ; \quad \underbrace{u \geq t+\Delta t}_{any\ time > t+\Delta t} \ , \tag{4}$$

$\hookrightarrow$ Optimal policy is not *time-consistent*.
The strategy which solves problem (3) has been called the *pre-commitment* policy (Basak,Chabakauri: 2010; Bjork et al: 2010)

- Much discussion on the economic meaning of such strategies.
- Possible to formulate a time-consistent version of mean-variance.
- Different applications may require different strategies.

Ulysses and the Sirens: A pre-commitment strategy

Ulysses had himself tied to the mast of his ship (and put wax in his sailor's ears) so that he could hear the sirens song, but not jump to his death.

Pre-commitment

Problem:

- Since the pre-commitment strategy is not time consistent, there is no natural dynamic programming principle
- We would like to formulate this problem as the solution of an HJB equation.
- How are we going to do this?

Solution:

- Use embedding technique [3] (Zhou and Li (2000), Li and Ng (2000))

[3] Does not require that we have convex constraints. Can be applied to problems with nonlinear transaction costs. Contrast with Lagrange multiplier approach.

Embedding

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

$$\sup_{c(X(u),u\geq t),c(\cdot)\in\mathbb{Z}} \Big\{ \underbrace{E^c_{t,x}[W_a(T)]}_{Reward} - \lambda \underbrace{Var^c_{t,x}[W_a(T)]}_{Risk} \Big\} ,$$

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

$\rightarrow \exists\ \gamma$ such that $c^*(\cdot)$ minimizes

$$\inf_{c(\cdot)\in\mathbb{Z}} E^{c(\cdot)}_{t,x}\left[\left(W_a(T) - \frac{\gamma}{2}\right)^2\right] . \tag{6}$$

Note we have effectively replaced parameter λ by γ in (6).

Construction of Efficient Frontier

We can alternatively now regard γ as a parameter, and determine the optimal strategy $c^*(\cdot)$ which solves

$$\inf_{c(\cdot)\in\mathbb{Z}} E_{t,x}^{c(\cdot)}\left[(W_a(T) - \frac{\gamma}{2})^2\right]$$

Once $c^*(\cdot)$ is known

- Easy to determine $E_{t,x}^{c^*(\cdot)}[W_a(T)]$, $Var_{t,x}^{c^*(\cdot)}[W_a(T)]$
- Repeat for different γ, traces out efficient frontier[4]

[4] Strictly speaking, since some values of γ may not represent points on the original frontier, we need to construct the upper left convex hull of these points (Tse, Forsyth, Li (2014), SIAM J. Control Optimization) .

Equivalence of MV optimization and target problem

MV optimization is equivalent[5] to investing strategy which

- Attempts to hit a target final wealth of $\gamma/2$
- There is a quadratic penalty for not hitting this wealth target
- From (Li and Ng(2000))

$$\underbrace{\frac{\gamma}{2}}_{wealth\ \ target} = \underbrace{\frac{1}{2\lambda}}_{risk\ \ aversion} + \underbrace{E^{c(\cdot)}_{t=0,x_0}[W_a(T)]}_{expected\ \ wealth}$$

Intuition: if you want to achieve $E[W_a(T)]$, you must aim higher

[5] Vigna, Quantitative Finance, to appear, 2014

HJB PIDE

Determination of the optimal control $c(\cdot)$ is equivalent to determining the value function

$$V(x,t) = \inf_{c \in \mathcal{Z}} \left\{ E_c^{x,t}[(W_a(T) - \gamma/2)^2] \right\} ,$$

Define:

$$\begin{aligned}
\mathcal{L}V &\equiv \frac{\sigma^2 s^2}{2} V_{ss} + (\mu - \lambda\kappa) s V_s + rbV_b - \lambda V \ , \\
\mathcal{J}V &\equiv \int_0^\infty p(\xi) V(\xi s, b, \tau) \ d\xi \\
& \quad p(\xi) = \text{ jump size density}
\end{aligned}$$

and the intervention operator $\mathcal{M}(c)\ V(s,b,t)$

$$\mathcal{M}(c)\ V(s,b,t) \quad = \quad V(S^+(s,b,c), B^+(s,b,c), t)$$

HJB PIDE II

The value function (and the control $c(\cdot)$) is given by solving the impulse control HJB equation

$$\max\left[V_t + \mathcal{L}V + \mathcal{J}V, V - \inf_{c \in \mathcal{Z}} (\mathcal{M}(c)\ V)\right] = 0$$
$$\text{if } (s + b > 0) \qquad (7)$$

Along with liquidation constraint if insolvent

$$V(s, b, t) = V(0, W_a(s, b), t)$$
$$\text{if } (s + b) \leq 0 \text{ and } s \neq 0 \qquad (8)$$

We can easily generalize the above equation to handle the discrete rebalancing case.

Computational Domain[6]

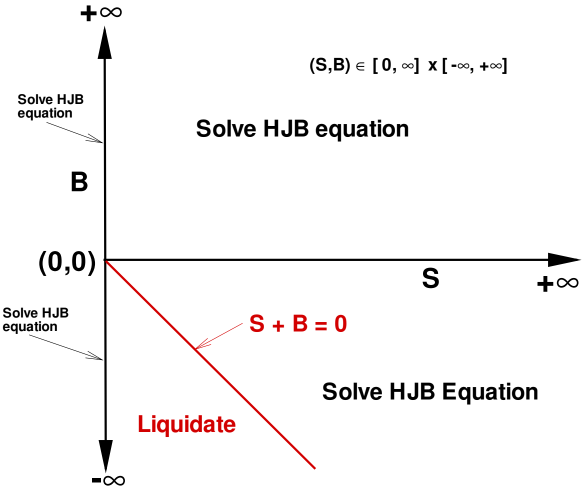

[6] If $\mu > r$ it is never optimal to short S

Well behaved utility function

Definition (Well-behaved utility functions)

A utility function $Y(W)$ is a well-behaved function of wealth W if it is an increasing function of W.

Proposition

Pre-commitment MV portfolio optimization is equivalent to maximizing the expectation of a well-behaved quadratic utility function if

$$W_a(T) \leq \frac{\gamma}{2}. \tag{9}$$

Obvious, since value function $V(x,t)$ is

$$V(x,t) = \sup_{c \in \mathcal{Z}} \left\{ E_c^{x,t}[Y(W_a(T))] \right\}$$

$$Y(W) = -(W - \gamma/2)^2$$

Dynamic MV Optimal Strategy

Theorem (Vigna (2014))

Assuming that (i) the risky asset follows a pure diffusion (no jumps), (ii) continuous re-balancing, (iii) infinite leverage permitted, (iv) trading continues even if bankrupt: then the optimal self-financing MV wealth satisfies

$$
\begin{aligned}
W_a(t) &\leq F(t) \; ; \; \forall t \\
& F(t) = \frac{\gamma}{2} e^{-r(T-t)} = \textit{ discounted wealth target}
\end{aligned}
$$

$\hookrightarrow$In this case, MV optimization

$\Rightarrow$ maximizes a well behaved quadratic utility function

Result can be generalized[7] to the case of

- Realistic constraints: finite leverage and no trading if insolvent
- But, we must have continuous rebalancing and no jumps

[7] Dang and Forsyth (2013)

Global Optimal Point

Examination of the HJB equation allows us to prove the following result (Dang and Forsyth (2013))

Lemma

The value function $V(s,b,t)$ is identically zero at

$$V(0, F(t), t) \ \equiv \ 0 \ ; \ F(t) = \frac{\gamma}{2} e^{-r(T-t)} \ , \ \forall t$$

Since $V(s,b,t) \geq 0$

- $V(0, F(t), t) = 0$ is a global minimum
- Any admissible policy which allows moving to this point is an optimal policy
- Once this point is attained, it is optimal to remain at this point

Globally Optimal Point[8]

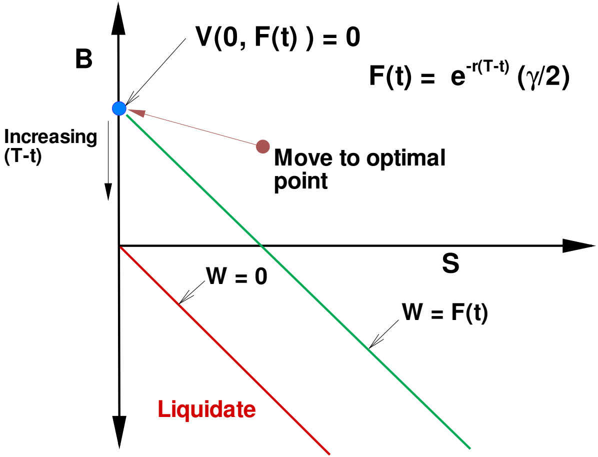

[8]This is admissible only if $\gamma > 0$

Optimal semi-self-financing strategy

Theorem (Dang and Forsyth (2013))

If $W_a(t) > F(t)$,[9] $t \in [0, T]$, *an optimal MV strategy is*[10]

- *Withdraw cash* $W_a(t) - F(t)$ *from the portfolio*
- *Invest the remaining amount* $F(t)$ *in the risk-free asset.*

Corollary (Well behaved utility function)

In the case of discrete rebalancing, and/or jumps, the optimal semi-self-financing MV strategy is

- *Equivalent to maximizing a well behaved quadratic utility function*[11]

---

[9] $F(t)$ is the discounted wealth target

[10] A similar semi-self-financing strategy for the discrete rebalancing case was first suggested in (Cui, Li, Wang, Zhu (2012) *Mathematical Finance*).

[11] Related to the idea of *time consistency in efficiency* from Li, Cui, Zhu (2011))

Intuition: Multi-period mean-variance

Optimal target strategy: try to hit $W_a(T) = \gamma/2 = F(T)$.

If $W_a(t) > F(t) = F(T)e^{-r(T-t)}$, then the target can be hit exactly by

- Withdrawing[12] $W_a(t) - F(t)$ from the portfolio
- Investing $F(t)$ in the risk free account

Optimal control for the target problem $\equiv$ optimal control for the Mean Variance problem

This strategy dominates any other MV strategy

$\rightarrow$ And the investor receives a bonus in terms of a free cash flow

[12] Idea that withdrawing cash may be mean variance optimal was also suggested in (Ehrbar, J. Econ. Theory (1990))

Numerical Method

We solve the HJB impulse control problem numerically using a finite difference method

- We use a semi-Lagrangian timestepping method
- Can impose realistic constraints on the strategy
  - Maximum leverage, no trading if insolvent
  - Arbitrarily shaped solvency boundaries
- Continuous or discrete rebalancing
- Nonlinearities
  - Different interest rates for borrowing/lending
  - Transaction costs
- Regime switching (i.e. stochastic volatility and interest rates)

We can prove[13] that the method is monotone, consistent, ℓ_∞ stable

$\rightarrow$ Guarantees convergence to the viscosity solution

[13] Dang and Forsyth (2014) Numerical Methods for PDEs

Numerical Examples

initial allocated wealth ($W_a(0)$)	100
r (risk-free interest rate)	0.04450
T (investment horizon)	20 (years)
$q_{\max}$ (leverage constraint)	1.5
$t_{i+1} - t_i$ (discrete re-balancing time period)	1.0 (years)

	mean downward jumps	mean upward jumps
μ (drift)	0.07955	0.12168
λ (jump intensity)	0.05851	0.05851
σ (volatility)	0.17650	0.17650
mean log jump size	-0.78832	0.10000
compensated drift	0.10862	0.10862

Efficient Frontier: continuous rebalancing

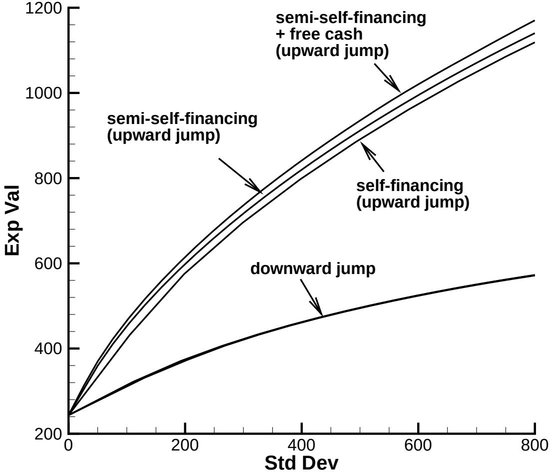

Figure: $T = 20$ years, $W_a(0) = 100$.

Efficient Frontier: discrete rebalancing

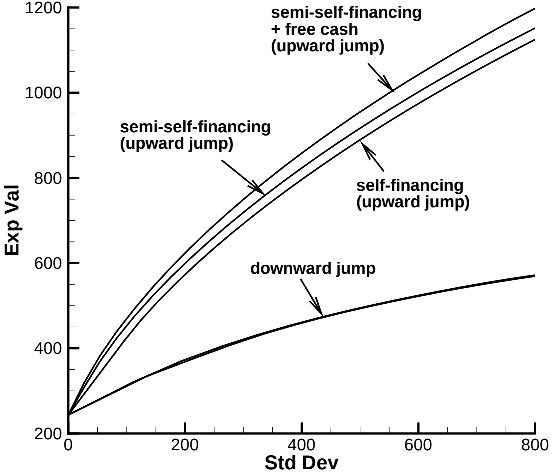

Figure: $T = 20$ years, $W_a(0) = 100$.

Efficient Frontier: zoom of downward jumps

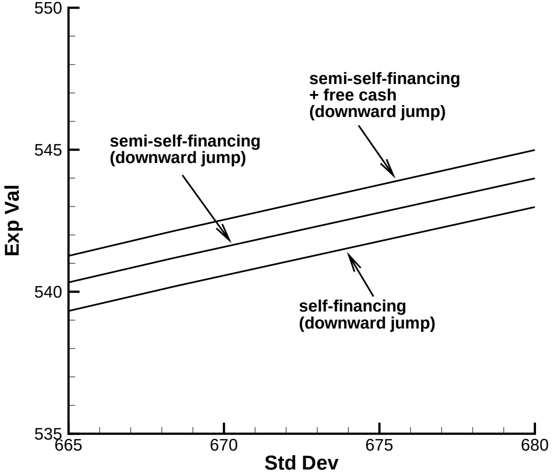

Figure: $T = 20$ years, $W_a(0) = 100$. Discrete rebalancing.

Example II

Two assets: risk-free bond, index

- Risky asset follows GBM (no jumps)

According to Benjamin Graham[14], most investors should

- Pick a fraction p of wealth to invest in an index fund (i.e. $p = 1/2$).
- Invest $(1 - p)$ in bonds
- Rebalance to maintain this asset mix

How much better is the optimal asset allocation vs. simple rebalancing rules?

[14] Benjamin Graham, *The Intelligent Investor*

Long term investment asset allocation

Investment horizon (years)	30
Drift rate risky asset μ	.10
Volatility σ	.15
Risk free rate r	.04
Initial investment W_0	100

Benjamin Graham strategy

Constant proportion	Expected Value	Standard Deviation	Quantile
$p = 0.0$	332.01	NA	NA
$p = 0.5$	816.62	350.12	$Prob(W(T) < 800) = 0.56$
$p = 1.0$	2008.55	1972.10	$Prob(W(T) < 2000) = 0.66$

Table: Constant fixed proportion strategy. $p =$ fraction of wealth in risky asset. Continuous rebalancing.

Optimal semi-self-financing asset allocation

Fix expected value to be the same as for constant proportion $p = 0.5$.

Determine optimal strategy which minimizes the variance for this expected value.

- We do this by determining the value of $\gamma/2$ (the wealth target) by Newton iteration

Strategy	Expected Value	Standard Deviation	Quantile
Graham $p = 0.5$	816.62	350.12	$Prob(W(T) < 800) = 0.56$
Semi-self-financing	816.62	142.85	$Prob(W(T) < 800) = 0.19$

Table: $T = 30$ years. $W(0) = 100$. Semi-self-financing: no trading if insolvent; maximum leverage $= 1.5$, rebalancing once/year.

Standard deviation reduced by 250 %, shortfall probability reduced by $3\times$

Cumulative Distribution Functions

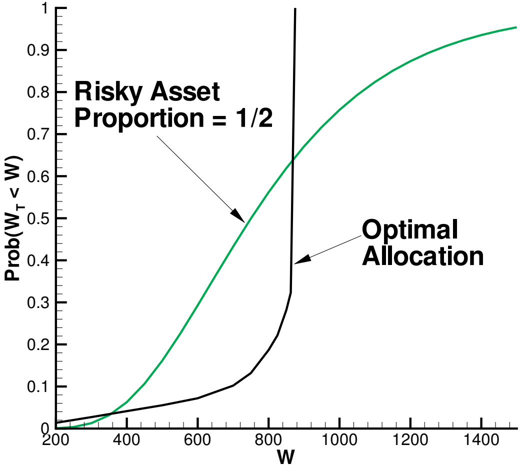

$E[W_T] = 816.62$ for both strategies

Optimal policy: $\uparrow W$ risk off; $\downarrow W(t)$ risk on

Optimal allocation gives up gains $\gg$ target in order to reduce variance and probability of shortfall.

Investor must pre-commit to target wealth

Why is the optimal semi-self-financing strategy so good?

Strategy	Expected Value	Standard Deviation	Quantile
$p = 0.5$	816.62	350.12	$Prob(W(T) < 800) = 0.56$
Optimal	816.62	142.85	$Prob(W(T) < 800) = 0.19$

Recall pre-commitment mean variance $\equiv$ quadratic target

- The investor pre-commits to a target final wealth
- Lower standard deviation, lower probability of shortfall
  - $\rightarrow$ The investor gives up some large possible gains $\gg$ expected value
- The investor commits to a long-term investment horizon
  - $\rightarrow$ Result not so impressive if $T < 20$ years
- The investor is not too greedy
  - $\rightarrow$ The target wealth $\gamma/2$ should be not too much larger than the expected value

$$\underbrace{\frac{\gamma}{2}}_{wealth\ \ target} = \underbrace{\frac{1}{2\lambda}}_{risk\ \ aversion} + \underbrace{E^{c(\cdot)}_{t=0,x_0}[W_a(T)]}_{expected\ \ wealth}$$

Open Question: Is it never optimal to withdraw if $W(t) < F(t)$?

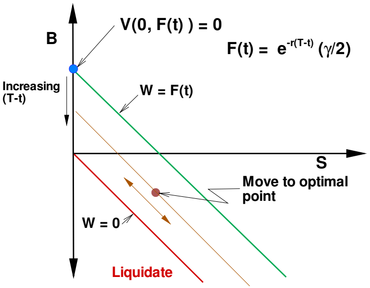

Seems obvious financially

Can prove this for jumps + continuous rebalancing

Proof does not work for jumps + discrete rebalancing

Problem with jumps and discrete rebalancing?

- Can jump into a region with large leverage, if another jump occurs $\rightarrow$ big loss

Numerical experiments

$\rightarrow$ Never optimal to withdraw if $W(t) < F(t)$

Conclusions

- Pre-commitment mean variance strategy
  - Equivalent to quadratic target strategy
- Semi-self-financing, pre-commitment mean variance strategy
  - Minimizes quadratic loss w.r.t. a target
  - Dominates self-financing strategy
  - Extra bonus of free cash-flow
- Example: target date fund
  - Optimal strategy dominates simple constant proportion strategy by a large margin
    - $\rightarrow$ Probability of shortfall $\simeq$ 3 times smaller!
  - But
    - $\rightarrow$ Investors must pre-commit to a wealth target
- Optimal stochastic control: teaches us an important life lesson
  - Decide on a life target ahead of time and stick with it
  - If you achieve your target, do not be greedy and want more