Optimal Trade Execution: Mean Variance or Mean Quadratic Variation?

Peter Forsyth[1] Stephen Tse[1] H. Windcliff[2] S. Kennedy[2]

[1]Cheriton School of Computer Science
University of Waterloo

[2]Morgan Stanley
New York

Oxford, February 24, 2012

The Basic Problem

Broker buys/sells large block of shares on behalf of client

- Large orders will incur costs, due to price impact (liquidity) effects
  - $\rightarrow$ e.g. rapidly selling a large block of shares will depress the price
- Slow trading minimizes price impact, but leaves exposure to stochastic price changes
- Fast trading will minimize risk due to random stock price movements, but price impact will be large
- What is the optimal strategy?

An Interesting Example of Price Impact

Remember Jerome Kerviel

- Rogue trader at Societe Generale
- The book value of Kerviel's portfolio, January 19, 2008 [1]
  $\rightarrow$ -2.7 Billion €
- SocGen decided to unwind this portfolio as rapidly as possible
- Over three days, the total cost of unwinding the portfolio was
  $\rightarrow$ -6.3 Billion €
- The price impact of rapid liquidation caused the realized loss to more than double the book value loss

[1] Report of the Commission Bancaire

Previous Approaches (a small sample)

Almgren, Chriss (2001) Mean-variance trade-off, discrete time, assume optimal asset positions are asset price independent (industry standard approach)

He, Mamaysky; Vath, Mnif, Pham; Schied, Schoneborn Maximize utility function, continuous time, dynamic programming, HJB equation.

Almgren, Lorenz (2011) Recognize that asset price independent solution is not optimal. Suggest HJB equation, continuous time, mean variance tradeoff.

Guilbaud, Mnif, H. Pham Impulse control formulation (discrete trading)

Gatheral, Schied (2011) Almgren, Chriss + GBM

Formulation

$$
\begin{aligned}
P &= \text{Trading portfolio} \\
  &= B + AS \\
B &= \text{Bank account: keeps track of gains/losses} \\
S &= \text{Price of risky asset} \\
A &= \text{Number of units of the risky asset} \\
T &= \text{Trading horizon}
\end{aligned}
$$

For Simplicity: Sell Case Only

Sell

$$t = 0 \to B = 0, S = S_0, A = A_0$$
$$t = T \to B = B_L, S = S_T, A = 0$$

- B_L is the cash generated by trading in $[0, T)$
  - $\hookrightarrow$ Plus a final sale at $t = T$ to ensure that zero shares owned.
- Success is measured by B_L (proceeds from sale).
- Maximize $E[B_L]$, minimize $Var[B_L]$

Price Impact Modelling

In practice, a hierarchy of models is used

Level 1 Considers all buy/sell orders of a large financial institution, over many assets

- Simple model of asset price movements, considers correlation between assets
- Output: *"sell* 10^7 *shares of RIM today"*.

Level 2 Single name sell strategy (schedule over the day)

- Level 2 models attempt to determine optimal strategy for selling a single name, assuming trades occur continuously, at rate v
- Price impact is a function of trade rate
- Output: *"sell* 10^5 *shares of RIM between 10:15-10:45"*

Level 3 Fine grain model

- Level 3 models assume discrete trades, and try to trade optimally based on an order book model.
- Output: *"place sell order for* 1000 *shares at 10:22"*

We focus on Level 2 models today.

Basic Problem

Trading rate v ($A =$ number of shares)

$$\frac{dA}{dt} = v \ .$$

Suppose that S follows geometric Brownian Motion (GBM) under the objective measure

$$\begin{aligned} dS &= (\eta + g(v))S\ dt + \sigma S\ dZ \\ &\quad \eta \text{ is the drift rate of } S \\ &\quad g(v) \text{ is the permanent price impact} \\ &\quad \sigma \text{ is the volatility} \\ &\quad dZ \text{ is the increment of a Wiener process} \quad . \end{aligned}$$

Basic Problem II

To avoid round-trip arbitrage (Huberman, Stanzl (2004))

$$g(v) = \kappa_p v$$

κ_p permanent price impact factor (const.)

The bank account B is assumed to follow

$$\frac{dB}{dt} = rB\ dt - vS_{exec}$$

r is the risk-free return

S_{exec} is the execution price

$= Sf(v)$

$f(v)$ is the temporary price impact

$(-vS_{exec})$ represents the rate of cash generated when buying shares at price S_{exec} at rate v.

Temporary Price Impact

Temporary price impact and transaction cost function $f(v)$ is assumed to be

$$
\begin{aligned}
f(v) \quad &= \quad [1 + \kappa_s \operatorname{sgn}(v)] \exp[\kappa_t \operatorname{sgn}(v)|v|^\beta] \\
&\quad\quad \kappa_s \text{ is the bid-ask spread parameter} \\
&\quad\quad \kappa_t \text{ is the temporary price impact factor} \\
&\quad\quad \beta \text{ is the price impact exponent}
\end{aligned}
$$

$$
\begin{aligned}
f(v) \quad &> 1 \qquad \text{if buying: execution price} > \text{pre-trade price} \\
&< 1 \qquad \text{if selling: execution price} < \text{pre-trade price}
\end{aligned}
$$

Optimal Strategy

Define:

$$
\begin{aligned}
X &= (S(t), A(t), B(t)) \;=\; \text{State} \\
B_L &= \text{Liquidation Value} \\
v(X, t) &= \text{trading rate}
\end{aligned}
$$

Let

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

so that

$$
\mathrm{Var}_{t,x}^{v(\cdot)}[B_L] = E_{t,x}^{v(\cdot)}[(B_L)^2] - \left(E_{t,x}^{v(\cdot)}[B_L] \right)^2
$$

Liquidation Value

- If $(S(T^-), A(T^-), B(T^-))$ are the state variables the instant before the end of trading $t = T^-$, B_L is given by

$$
\begin{aligned}
B_L &= B - v_T (\Delta t)_T S f(v_T) \\
& v_T = \frac{0 - A}{(\Delta t)_T}
\end{aligned}
$$

.

- Choosing $(\Delta t)_T$ small, penalizes trader for not hitting target $A = 0$.
- Optimal strategy will avoid the state $A \neq 0$
- Numerical solution insensitive to $(\Delta t)_T$ if sufficiently small

Mean Variance: Standard Formulation

The objective is to determine the strategy $v(\cdot)$ such that

$$
\begin{aligned}
J(x,t) \;&=\; \sup_{v(X(u),u\geq t)} \left\{ E^{v}_{t,x}[B_L] - \lambda Var^{v}_{t,x}[B_L] \right\} , \\
&\lambda = \text{ Lagrange multiplier}
\end{aligned} \qquad (1)
$$

Solving (1) for various λ traces out a curve in the expected value, standard deviation plane.

• Let $v^*_t(x,u), u \geq t$ be the optimal policy for (1).
Then $v^*_{t+\Delta t}(x,u), u \geq t+\Delta t$ is the optimal policy for

$$
\begin{aligned}
&J(X(t+\Delta t), t+\Delta t) = \\
&\quad \sup_{v(X(u),u\geq t+\Delta t))} \left\{ E^{v}_{t+\Delta t, X(t+\Delta t)}[B_L] - \lambda Var^{v}_{t+\Delta t, X(t+\Delta t)}[B_L] \right\} .
\end{aligned}
$$

Pre-commitment Policy

However, in general

$$v_t^*(X(u), u) \neq v_{t+\Delta t}^*(X(u), u) \; ; \; u \geq t + \Delta t \; , \qquad (2)$$

$\hookrightarrow$ Optimal policy is not *time-consistent*.
The strategy which solves problem (1) 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.
- Or other strategies: mean quadratic variation
- Different applications may require different strategies.
- We focus on pre-commitment solution today, with a brief discussion of alternative 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:

- Go back to first principles

Minimum Variance: Basic Principle

Equivalent formulation: determine the strategy $v(\cdot)$ such that

$$\min Var_{t,x}^{v(\cdot)}[B_L] = E_{t,x}^{v(\cdot)}[(B_L)^2] - d^2$$

$$\text{subject to } \begin{cases} E_{t,x}^{v(\cdot)}[B_L] = d \\ v(\cdot) \in \mathbb{Z} \end{cases}$$

$$\mathbb{Z} = \text{ set of admissible controls}$$

Given an expected return $d = E_{t,x}^{v(\cdot)}[B_L]$, strategy $v(\cdot)$ produces the smallest possible variance.

Varying the parameter d traces out a curve in the expected value - standard deviation plane.

Eliminate Constraint

Original problem is convex optimization, use Lagrange multiplier γ to eliminate constraint.

$$\max_{\gamma} \min_{v(\cdot)\in\mathbb{Z}} E_{t,x}^{v(\cdot)}\left[(B_L)^2 - d^2 - \gamma(E_{t,x}^{v(\cdot)}[B_L] - d)\right] . \tag{3}$$

Suppose somehow we know the γ which solves (3), for fixed d.

Then the optimal strategy $v^*(\cdot)$ which solves (3) is given by (for fixed γ)

$$\min_{v(\cdot)\in\mathbb{Z}} E_{t,x}^{v(\cdot)}\left[\left(B_L - \frac{\gamma(t,x,d)}{2}\right)^2\right] . \tag{4}$$

Note we have effectively replaced parameter d by γ in (4).

Construction of Efficient Frontier [2]

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

$$\min_{v(\cdot)\in\mathbb{Z}} E_{t,x}^{v(\cdot)}\left[(B_L - \frac{\gamma}{2})^2\right] . \tag{5}$$

Once $v^*(\cdot)$ is known, we can easily determine $E_{t,x}^{v^*(\cdot)}[B_L]$, $E_{t,x}^{v^*(\cdot)}[(B_L)^2]$, by solving an additional linear PDE.

For given γ, this gives us $(E_{t,x}^{v^*(\cdot)}[B_L], \mathrm{Std}_{t,x}^{v^*(\cdot)}[B_L])$, a single point on the efficient frontier.

Repeating the above for different γ generates points on the efficient frontier.
$\hookrightarrow$ Efficient frontier construction reduces to repeated solves of (5).

[2] In simple cases, this can be shown to be equivalent to minimizing quadratic loss, with target $\simeq$ expected value (Vigna, 2011)

Hamilton Jacobi Bellman (HJB) Equation

Let

$$
\begin{aligned}
& V(s, \alpha, b, \tau) \\
& = \min_{v(\cdot) \in \mathbb{Z}} \left\{ E_{t,x}^{v(\cdot)} \left[(B_L - \frac{\gamma}{2})^2 \;\middle|\; S(t) = s, A(t) = \alpha, B(t) = b \right\} \right. \\
& \quad x = (s, \alpha, b) \\
& \quad s = \text{ stock price} \\
& \quad \alpha = \text{ number of units of stock} \\
& \quad b = \text{ cash obtained so far} \\
& \quad T = \text{ Trading horizon} \\
& \quad \tau = T - t \\
& \quad \mathbb{Z} = [v_{\min}, 0] \qquad \text{(Only selling permitted)}
\end{aligned}
$$

HJB Equation for Optimal Control $v^*(\cdot)$

We can use dynamic programming[3] to solve for

$$\min_{v(\cdot)\in\mathbb{Z}} E_{t,x}^{v(\cdot)}\left[(B_L - \frac{\gamma}{2})^2\right] . \tag{6}$$

Then, using usual arguments, $V(s,\alpha,b,\tau)$ is determined by

$$\begin{aligned} V_\tau &= \mathcal{L}V + rbV_b + \min_{v\in\mathbb{Z}}\left[-vsf(v)V_b + vV_\alpha + g(v)sV_s\right] \\ &\mathcal{L}V \equiv \frac{\sigma^2 s^2}{2}V_{ss} + \eta s V_s \\ &\mathbb{Z} = [v_{min}, 0] \end{aligned}$$

with the payoff $V(s,\alpha,b,\tau=0) = (b_L^2 - \gamma/2)$. [4]

[3] But this is not time-consistent since $\gamma = \gamma(t,x,d)$

[4] But note that v is arbitrary if $V_b = V_\alpha = V_s = 0$

But solving the HJB equation requires some work

I will give a brief description of how to do this (later).

- But this is considered too complex in industry
- So, the original (Almgren and Chriss) paper made several approximations (e.g. $v(\cdot)$ independent of $S(t)$).
- In fact, a careful read of this paper, shows that the objective function (after the approximations) is not actually mean-variance, but is mean quadratic-variation

Mean Quadratic Variation

Formally, the quadratic variation risk measure is defined as

$$E\left[\int_t^T \big(A(t')dS(t')\big)^2\right]. \tag{7}$$

Informally (if $P = B + AS$)

$$\big(A(t')dS(t')\big)^2 = \big(dP(t')\big)^2$$

i.e. the quadratic variation of the portfolio value process.

Originally suggested as an alternate risk measure by Brugierre (1996).

This measures risk in terms of the variability of the stock holding position, along the entire trading path.

Mean Quadratic Variation

Objective Function

$$J(s,\alpha,t,v(\cdot);\lambda) = E^{v(\cdot)}_{s,\alpha,t}\Big[B_L\Big] - \lambda E^{v(\cdot)}_{s,\alpha,t}\Big[\int_t^T \big(A(t')dS(t')\big)^2\Big] \quad (8)$$

where

$$B_L = \int_t^{T^-} (\mathit{Cash\ Flows\ from\ selling})dt' + (\mathit{Final\ Sale\ at\ } t = T) \quad (9)$$

One can easily derive the HJB equation for the optimal control $v^*(\cdot)$

$$\begin{aligned} V_\tau &= \eta s V_s + \frac{\sigma^2 s^2}{2} V_{ss} - \lambda\sigma^2\alpha^2 s^2 \\ &\qquad + \max_{v\in\mathbb{Z}}\Big[e^{r\tau}(-vf(v))s + g(v)sV_s + vV_\alpha\Big]\ . \end{aligned}$$

Mean Quadratic Variation

- The control is time consistent in this case
- If we assume Arithmetic Brownian Motion, then HJB equation has analytic solution (Almgren, Chriss(2001))
  - Control is independent of $S(t)$

One could argue that mean quadratic variation is a reasonable risk measure

- Risk is measured along the entire trading path
- In contrast, Mean variance only measures risk at end of path
- Time-consistency $\rightarrow$ smoothly varying controls

But

$$\text{Mean Quadratic Variation} \neq \text{Mean Variance}$$

How do We Measure Performance of Trading Algorithms?

Imagine we carry out many hundreds of trades

We then examine post-trade data[5]

- Determine the realized mean return and standard deviation
- Assuming the modeled dynamics very closely match the dynamics in the real world
  - → Optimal pre-commitment Mean Variance strategy will result in the largest realized mean return, for given standard deviation

So, if we measure performance in this way

- We should use Mean Variance optimal control
- But this is not what's done in industry
  - → Effectively, a Mean Quadratic Variation Control is used (Almgren, Chriss (2001))

---

[5] According to my industry contacts, some clients actually do this

HJB Equation: Mean Variance

Define the Lagrangian derivative

$$\frac{DV}{D\tau}(v) = V_\tau - V_s g(v)s - V_b(rb - vf(v)s) - V_\alpha v,$$

which is the rate of change of V along the characteristic curve

$$s = s(\tau)\ ;\ b = b(\tau)\ ;\ \alpha = \alpha(\tau)$$

defined by the trading velocity v through

$$\frac{ds}{d\tau} = -g(v)s, \quad \frac{db}{d\tau} = -(rb - vf(v)s), \quad \frac{d\alpha}{d\tau} = -v.$$

HJB Equation: Lagrangian Form

We can then write the Mean Variance HJB equation as

$$\mathcal{L}V - \max_{v(\cdot)\in Z} \frac{DV}{D\tau}(v) = 0.$$

$$\mathcal{L}V \equiv \frac{\sigma^2 s^2}{2} V_{ss} + \eta s V_s$$

Numerical Method:

- Discretize the Lagrangian form directly (semi-Lagrangian method)
- Timestepping algorithm
  - Solve local optimization problem at each grid node
  - Discretized linear PDE solve to advance one timestep
- Provably convergent to the viscosity solution of the HJB PDE
- Similar approach for the Mean Quadratic Variation HJB PDE

Numerical Method: Efficient Frontier

Recall that (Mean Variance)

$$V(s, \alpha, b, \tau = 0) \;\; = \;\; (b_L - \gamma/2)^2$$

Numerical Algorithm

- Pick a value for γ
  - Solve HJB equation for optimal control $v = v(s, \alpha, b, \tau)$
  - Store control at all grid points
  - Simulate trading strategy using a Monte Carlo method (use stored optimal controls)
  - Compute mean, standard deviation
  - This gives a single point on the efficient frontier
- Repeat

Similar approach for Mean Quadratic Variation

Numerical Examples

Simple case: GBM, zero drift, zero permanent price impact

$$dS = \sigma S \; dZ$$

Temporary Price Impact:

$$f(v) \quad = \quad exp(\kappa_t v)$$

T	r	s_{init}	α_{init}	Action	v_{min}
1/250 (One Day)	0.0	100	1.0	Sell	-1000/T

Case	σ	κ_t	Percentage of Daily Volume
1	1.0	2×10^{-6}	16.7%
2	0.2	2.4×10^{-6}	20.0%
3	0.2	6×10^{-7}	5.0%
4	0.2	1.2×10^{-7}	1.0%
5	0.2	2.4×10^{-8}	0.2%

$\sigma = 1.0$, 16.7% daily volume, $S_{init} = 100$

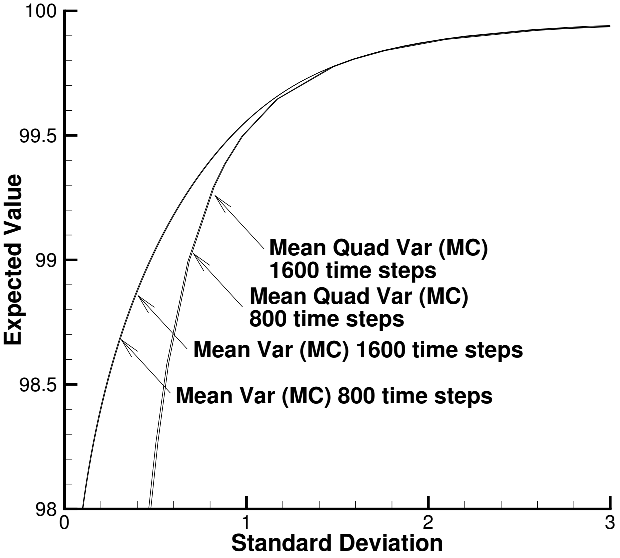

$\sigma = .2$, 20% daily volume, $S_{init} = 100$

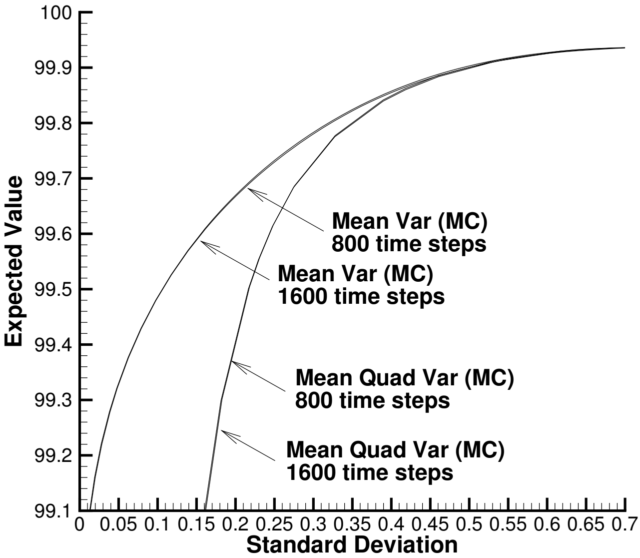

$\sigma = .2$, 5% daily volume, $S_{init} = 100$

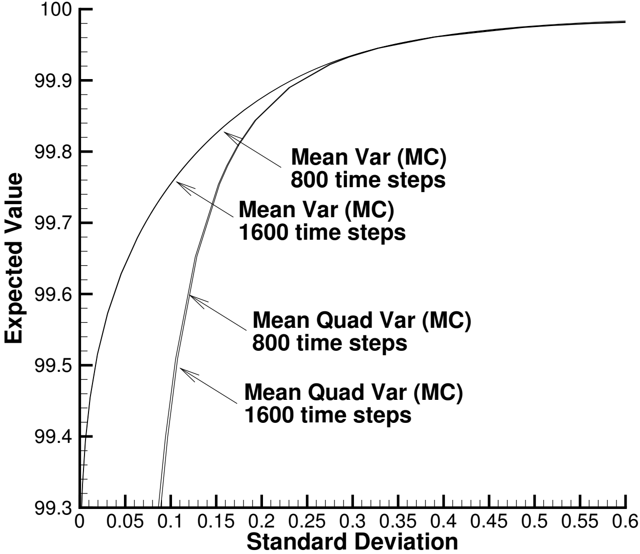

$\sigma = .2$, 1% daily volume, $S_{init} = 100$

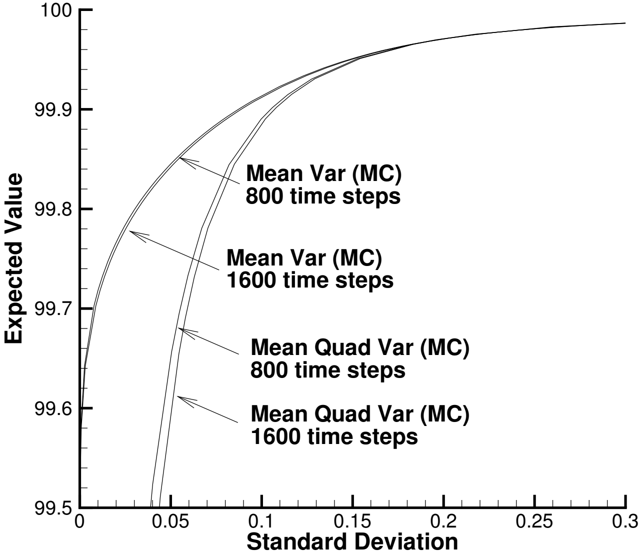

$\sigma = .2$, 0.2% daily volume, $S_{init} = 100$

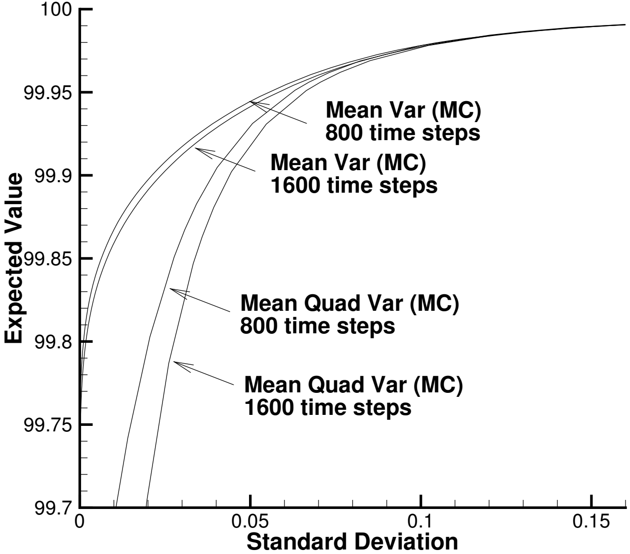

Optimal trading rate: $t = 0, \alpha = 1, b = 0$

- $\sigma = 1.0$, 16.7% daily volume
- Mean: 99.29.
- Std(Mean Variance) $= 0.68$
- Std(Mean Quadratic Variation) $= 0.93$
- $V_s \simeq V_b \simeq V_s \simeq 0$ when $S > 104$

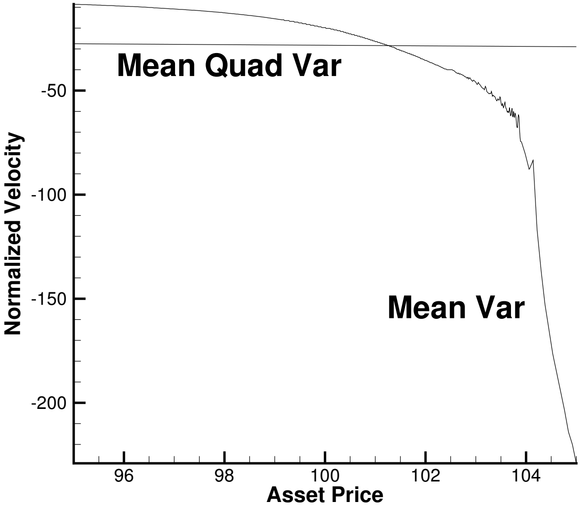

Mean Share Position (α) vs. Time

- $\sigma = 1.0$, 16.7% daily volume
- Mean: 99.29.
- Std(Mean Variance) $= 0.68$
- Std(Mean Quadratic Variation) $= 0.93$

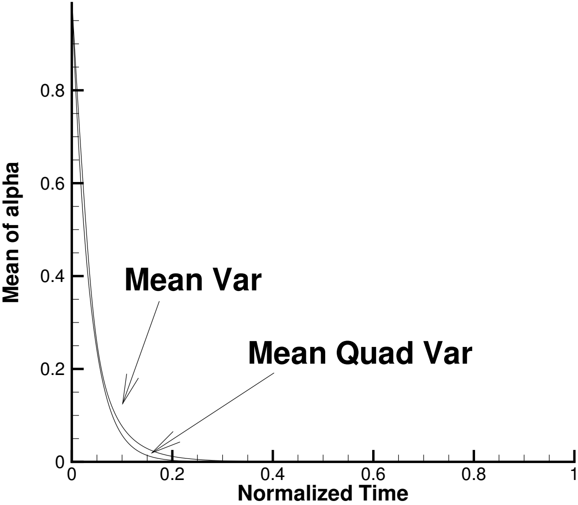

Standard Deviation of Share Position (α) vs. Time

- $\sigma = 1.0$, 16.7% daily volume
- Mean: 99.29.
- Std(Mean Variance) $= 0.68$
- Std(Mean Quadratic Variation) $= 0.93$

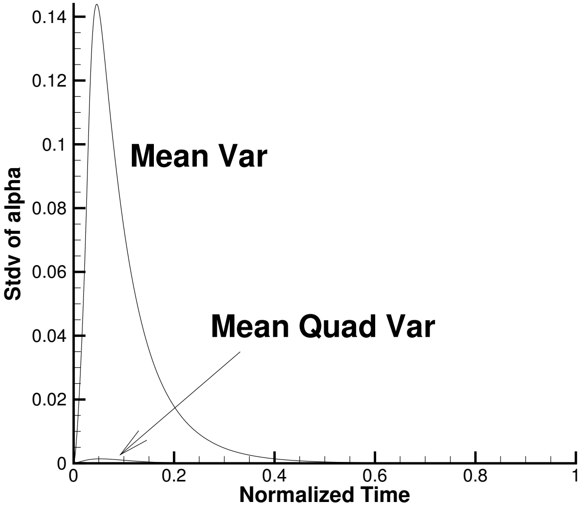

Conclusions: Mean Variance

Pros:

- If performance is measured by post-trade data (mean and variance)
  - $\rightarrow$ This is the truly optimal strategy
- Significantly outperforms Mean Quadratic Variation for low levels of required risk (fast trading)

Cons:

- Non-trivial to compute optimal strategy
- Very aggressive in-the-money strategy
- Share position has high standard deviation
- Optimal trading rate is *almost ill posed*: many nearby strategies give almost same efficient frontier
  - $\rightarrow$ Simple example: zero standard deviation [6]

[6] Recall that v is arbitrary if $V_s = V_b = V_\alpha = 0$

Conclusions: Mean Quadratic Variation

Pros:

- Simple analytic solution for Arithmetic Brownian Motion Case
- Trading rate a smooth, predictable function of time
  - $\rightarrow$ For GBM case, only weakly sensitive to asset price S
- Almost same results as Mean Variance, for large levels of required risk (slow trading)

Cons:

- If performance is measured by post-trade data (mean and variance)
  - $\rightarrow$ This is not the optimal strategy
  - $\rightarrow$ Significantly sub-optimal for low levels of risk (fast trading)