An HJB Equation Approach to Optimal Trade Execution

Peter Forsyth,
Cheriton School of Computer Science
University of Waterloo

Linz, November 19, 2008

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
- 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?

Previous Approaches

Almgren, Chriss Mean-variance trade-off, discrete time, assume optimal asset positions are path-independent

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

Almgren, Lorenz Recognize that path-independent solution is not optimal. Suggest HJB equation, continuous time, mean variance tradeoff. No results.

Lorenz (2008) Mean variance tradeoff: analytic solution for simple cases

Formulation

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

For Simplicity: Sell Case Only

Sell

$$t = 0 \rightarrow B = 0, S = S_0, \alpha = \alpha_{sell}$$

$$t = T \rightarrow B = B_L, S = S_T, \alpha = 0$$

- 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

Assume trades occur instantaneously in discrete amounts, leads to impulse control formulation.
$\hookrightarrow$ Problem: the price impact of two discrete trades independent of time interval between trades (unrealistic).

Alternate approach: assume trades occur continuously, at trading rate v.
$\hookrightarrow$ In this case, price impact can be a function of trade rate.
$\hookrightarrow$ Problem: real trading takes place discretely.

Neither model is perfect. We use the continuous trade model in the following.

Basic Problem

Trading rate v ($\alpha =$ number of shares)

$$\frac{d\alpha}{dt} = v \; .$$

Suppose that S follows geometric Brownian Motion (GBM)

$$dS = (\eta + g(v))S\ dt + \sigma S\ dZ$$

η is the drift rate of S

$g(v)$ is the permanent price impact

σ is the volatility

dZ is the increment of a Wiener process .

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

$$\frac{dB}{dt} = rB\ dt - vSf(v)$$

r is the risk-free return

$f(v)$ is the temporary price impact

$-vSf(v)$ represents the rate of cash generated when selling shares at price $Sf(v)$ at rate v.

Temporary Price Impact

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

$$f(v) = [1 + \kappa_s \,\mathrm{sgn}(v)] \exp[\kappa_t \,\mathrm{sgn}(v)|v|^{\beta}]$$

κ_s is the bid-ask spread parameter

κ_t is the temporary price impact factor

β is the price impact exponent

Control Problem

Select the control $v(t)$ (i.e. the selling strategy) so as to maximize

$$\max_{v(t)} \left(E^{t=0}[B_L] - \lambda Var^{t=0}[B_L] \right)$$
$$E^{t=0}[\cdot] = \text{ Expectation as seen at } t = 0$$
$$Var^{t=0}[\cdot] = \text{ Variance as seen at } t = 0 \qquad (1)$$

B_L is the total cash received from the selling strategy

Varying λ generates pairs $(E^{t=0}[B_L], Var^{t=0}[B_L])$ along the efficient frontier. More intuitive result than usual power-law/exponential utility function approach. (What is the utility function of a bank?)

The Liquidation Value

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

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

. • Choosing $(\Delta t)_T$ small, penalizes trader for not hitting target $\alpha = 0$.

• Optimal strategy will avoid the state $\alpha \neq 0$

$\hookrightarrow$ Numerical solution insensitive to $(\Delta t)_T$ if sufficiently small

Pre-committment vs. Time-consistent

We are maximizing (as seen at $t = 0$)

$$\max_{v(t)} \left(E^{t=0}[B_L] - \lambda Var^{t=0}[B_L] \right. \qquad (2)$$

This is the *pre-committment* policy, i.e. the strategy as a function of (S, B, α, t) is computed at $t = 0$.

• The trader follows this strategy even if (2) computed at $t > 0$ would yield a different strategy.

• The *pre-committment* policy is not *time-consistent* in this sense

• It is also possible to determine a time consistent mean-variance optimal strategy (Basak and Chabakauri, 2008)

How Do We Measure Success?

Suppose we lived in a world where our model of the process for the risky asset, and the price impact functions was perfect

- Suppose we followed the *pre-committment* policy for thousands of trades, of same stock

We then compute the mean and standard deviation of the trades

- Any other strategy (including time consistent) must result in a smaller mean gain for the same standard deviation, compared to the pre-committment strategy

$\hookrightarrow$ Time-consistent = pre-committment + constraints

Dynamic Programming and Efficient Frontier

We would like to use Dynamic Programming and derive an HJB equation for the optimal strategy $v^*(t)$.

$$\max_{v(t)} \left(E^{t=0}[B_L] - \lambda Var^{t=0}[B_L] \right)$$
$$E^{t=0}[\cdot] = \text{ Expectation}$$
$$Var^{t=0}[\cdot] = \text{ Variance} \quad (3)$$

But the variance term in the objective function causes difficulty. Solution (Li, Ng(2000); Zhou, Li (2000); theoretical analysis, not numerical)

Linear-Quadratic (LQ) Problem

Theorem 1 (Equivalent LQ problem). *If $v^*(t)$ is the optimal control of Mean-Variance problem (3) then $v^*(t)$ is also the optimal control of problem*

$$\max_{v(t)} E^{t=0}[\mu B_L - \lambda B_L^2] \tag{4}$$

where

$$\mu = 1 + 2\lambda E_{v^*}^{t=0}[B_L] \tag{5}$$

where v^* is the optimal control of problem (4).

LQ Problem II

At first glance, this does not seem to be very useful

• μ is a function of the optimal control v^*

$\hookrightarrow$ Not known until the problem is solved

Since $\lambda > 0$, we can rewrite the LQ problem as

$$\min_{v(t)} \left(E^{t=0}[(B_L - \frac{\gamma}{2})^2] - \frac{\gamma^2}{4} \right)$$

$$\gamma = \frac{\mu}{\lambda}$$

LQ Problem III

For fixed γ, an optimal control of the original problem is an optimal control of

$$\min_{v(t)} E^{t=0}[(B_L - \frac{\gamma}{2})^2] \quad . \tag{6}$$

Possible solution method: pick a value of γ, solve (6) for optimal strategy $v^*(t)$. Then, with known $v^*(t)$, compute

$$E_{v^*}^{t=0}[B_L] \quad ; \quad \lambda = \frac{1}{\gamma - 2E_{v^*}[B_L]}$$

Note: effectively parameter λ replaced by parameter γ.

Efficient Frontier

Choosing different values of γ
$\hookrightarrow$ Corresponds to different choices of λ
$\hookrightarrow$ Determines the optimal strategy $v_\lambda^*(t)$
$\hookrightarrow$ We generate pairs of points (for each λ)

$$\left(E_{v_\lambda^*}^{t=0}[B_L^2], E_{v_\lambda^*}^{t=0}[B_L]\right) \quad (7)$$

This can then be converted to points on the efficient frontier.

Varying $\gamma \rightarrow$ traces out efficient frontier.

A Better Method

Define a pseudo-bank account $\mathcal{B}(t)$

$$\mathcal{B}(t) = B(t) - \frac{\gamma e^{-r(T-t)}}{2} \quad . \tag{8}$$

so that the control problem becomes (γ disappears)

$$\min_{v(t)} E^{t=0}[\mathcal{B}_L^2] \quad . \tag{9}$$

Assume the state of the strategy is fully specified by the variables $(S, \mathcal{B}, \alpha, t)$.

A Better Method II

Let

$$V(S, \mathcal{B}, \alpha, t) \;\; = \;\; E^t_{v^*}[\mathcal{B}^2_L]$$

At $t = 0$ note that (assuming real bank account $= 0$)

$$\gamma = -2\mathcal{B}_0 e^{rT}$$

This means that if we examine $V(S_0, \mathcal{B}_0, \alpha_0, t = 0)$ for various $\mathcal{B}_0 \rightarrow$ we can determine $v^*_\lambda(t), E^{t=0}_{v^*_\lambda}[B^2_L])$ **for any** λ

Solution of the Optimal Control Problem

Recall $V = V(S, \mathcal{B}, \alpha, \tau = T - t) = E_{v^*}^{t=T-\tau}[\mathcal{B}_L^2]$. Let

$$\mathcal{L}V \equiv \frac{\sigma^2 S^2}{2} V_{SS} + \eta S V_S + r\mathcal{B}V_{\mathcal{B}} \quad .$$

Then, using usual arguments, $V(S, \mathcal{B}, \alpha, \tau)$ is determined by

$$\begin{aligned} V_\tau &= \mathcal{L}V + r\mathcal{B}V_{\mathcal{B}} + \min_{v \in Z}\left[-vSf(v)V_{\mathcal{B}} + vV_\alpha + g(v)SV_S \right] \\ & Z = [v_{min}, v_{max}] \end{aligned}$$

with the payoff $V(S, \mathcal{B}, \alpha, \tau = 0) = \mathcal{B}_L^2$.

Numerical Method

Step 1 Solve HJB equation once with initial condition $V(S, \mathcal{B}, \alpha, \tau = 0) = \mathcal{B}_L^2$.
$\longrightarrow$ This determines optimal control $v^*(t)$, $E_{v^*}^{t=0}[\mathcal{B}_L^2]$.

Step 2 Solve PDE problem again, using known control from Step 1, with initial condition $U(S, \mathcal{B}, \alpha, \tau = 0) = \mathcal{B}_L$, this gives $U = E_{v^*}^{t=0}[\mathcal{B}_L]$.

Step 3 Solution of these two PDEs allows us to generate points along **the entire efficient frontier**.

Numerical Method II

Nonlinear HJB equation solved using finite difference with semi-Lagrangian timestepping

- Optimal trade rate at each node determined by discretizing $[v_{\min}, v_{\max}]$, and using linear search (expensive but bullet proof)

- Consistent, stable, monotone $\longrightarrow$ converges to viscosity solution of HJB equation.

Further Simplification

Since $V = V(S, \mathcal{B}, \alpha, \tau)$, we need to solve a $3-d$ HJB equation.

For many reasonable price impact functions, we have the property ($V = E[\mathcal{B}_L^2], U = E[\mathcal{B}_L]$)

$$
\begin{aligned}
V(\xi S, \xi \mathcal{B}, \alpha, \tau) &= \xi^2 V(S, \mathcal{B}, \alpha, \tau) = \xi^2 E_{\alpha^*}^{t=0}[\mathcal{B}_L^2] \\
U(\xi S, \xi \mathcal{B}, \alpha, \tau) &= \xi U(S, \mathcal{B}, \alpha, \tau) = \xi E_{\alpha^*}^{t=0}[\mathcal{B}_L] \\
& \xi = \text{ constant}
\end{aligned}
$$

This similarity reduction can be used to reduce the $3-d$ PDEs to $2-d$ PDEs.

Optimal Liquidation Example

Parameter	Value
Volatility σ	.40
Trading Horizon (years) T	1/12
Drift Rate η	.10
Risk Free Rate r	0.0
Pretrade Price S_0	100
Initial Shares α_{sell}	1.0
Permanent Impact Factor κ_p	0.0
Temporary Impact Factor κ_t	.002
Relative Bid-Ask Spread κ_s	0.0
Temporary Impact Exponent β	1.0
Minimum Trading Rate v_{min}	-25/T
Maximum Trading Rate v_{max}	0.0

Efficient Frontier

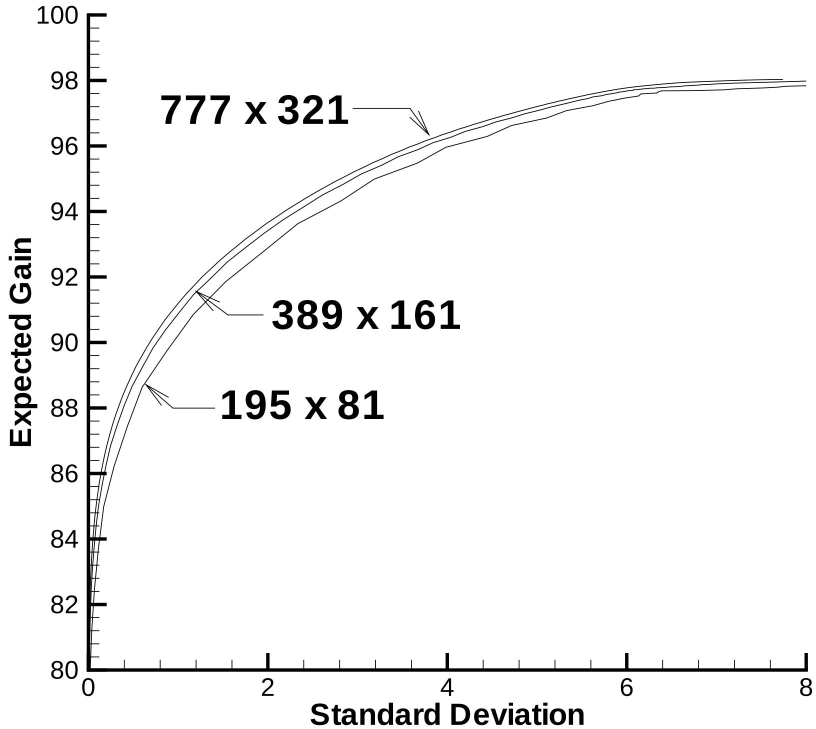

- $n_S \times n_\alpha = (S \text{ nodes}) \times (\alpha \text{ nodes})$, $\mathcal{B}$ eliminated by similarity reduction
- Average price obtained during one month period, pre-trade price 100

Sanity Check: Two Simple Cases

Trade at constant rate $v = 1/T = 12$. Minimizes price impact, (approximately) maximizes expected gain.

$$E[e^{-\kappa_t/T}\frac{1}{T}\int_0^T S(t)\ dt] \simeq 98.04$$

$$\text{Standard Deviation}[e^{-\kappa_t/T}\frac{1}{T}\int_0^T S(t)\ dt] \simeq 6.55\ .$$

Trade at maximum possible rate, minimizes standard deviation

- Expected gain $=$ standard deviation $= 0$.

Two Simple Cases II

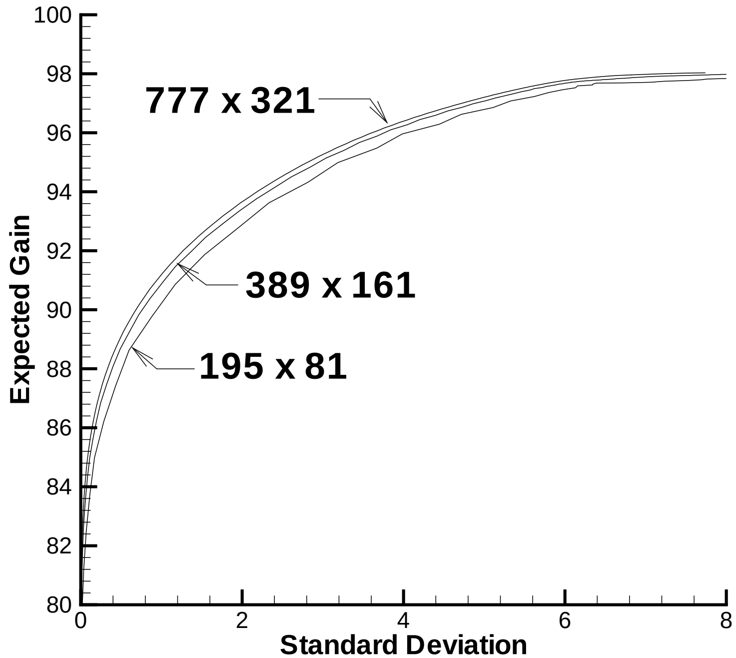

- point $(6.55, 98.04)$, maximizes gain
- point $(0, 0)$ minimizes standard deviation

Optimal strategy

Consider the point on the efficient frontier

$$\begin{aligned} \text{Expected Gain} &= 96.17 \\ \text{Standard Deviation} &= 3.76 \end{aligned}$$

• Examine the value of $v^*(S, B = 0, \alpha = 1, t = 0)$
• Recall that we assume $S = 100$ at $t = 0$
• This means that we can interpret this as examining how the optimal strategy would change if the asset price suddenly moved away from $S = 100$.

Optimal strategy: $v^*(S, B = 0, \alpha = 1, t = 0)$

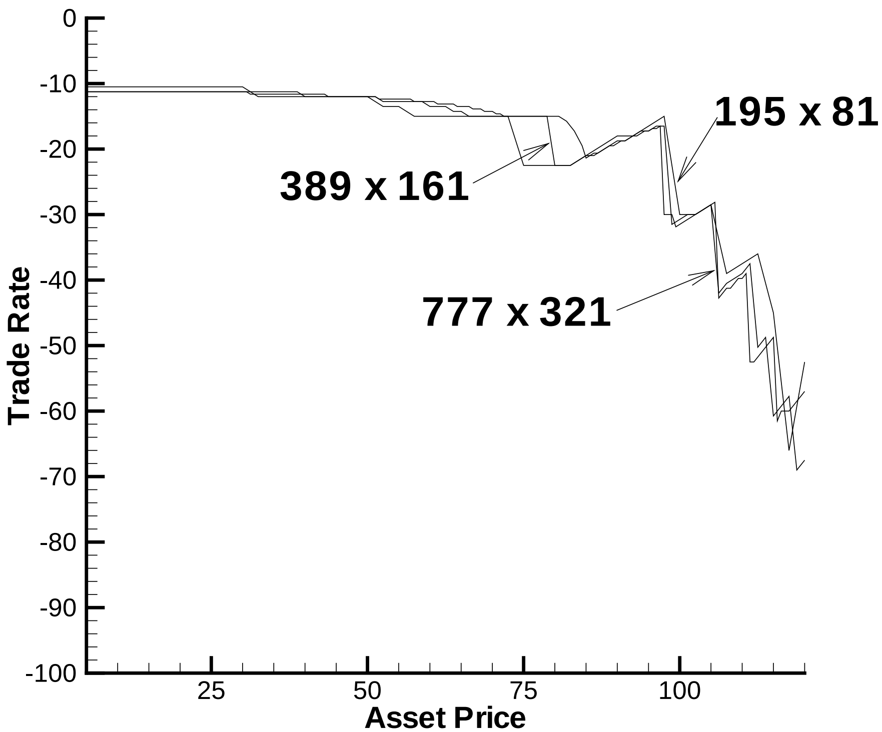

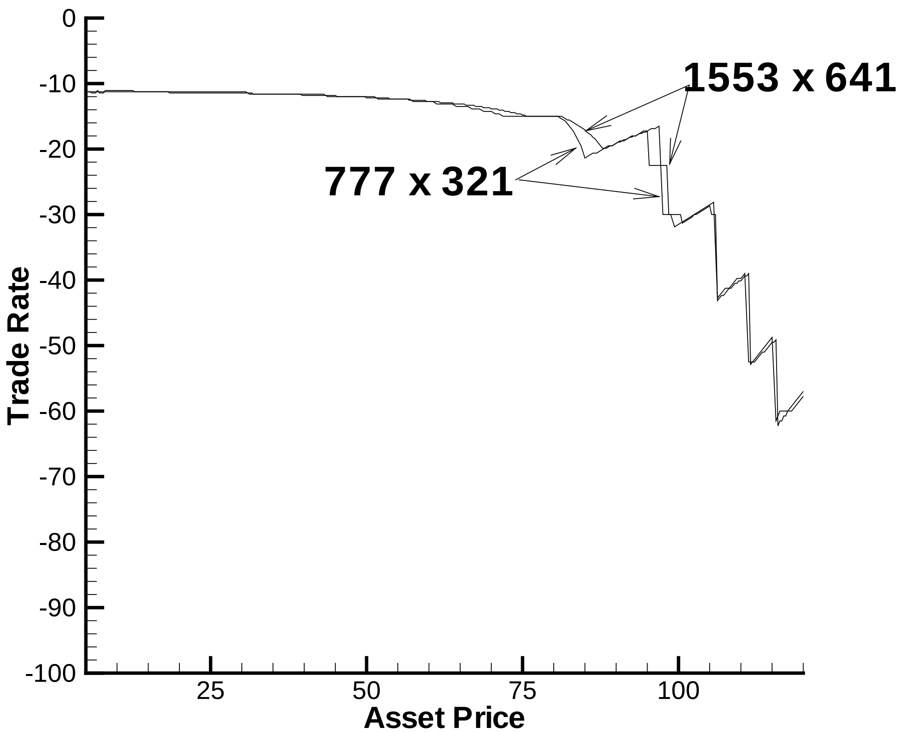

Sawtooth Pattern?

Why do we get this sawtooth pattern in the optimal trade rate as a function of S?

How can the optimal strategy be a non-smooth function of S?

Semi-Lagrangian timestepping

$\hookrightarrow$ Optimal trade at each point $(S, \mathcal{B}, \alpha, t)$ determined by finding the minimum value of the objective function along a curve in $(S, \mathcal{B}, \alpha, t)$ space.

• We plot the value of the objective function along this curve, for fixed initial point $(\mathcal{B}, \alpha)$ at two nearby initial points in the S direction.

Local Objective Function

- Optimal trade rate determined by global minimum

- Objective function has multiple local minima

- Objective function is very flat between local minima

- Location of global minimum is not a continuous function of S

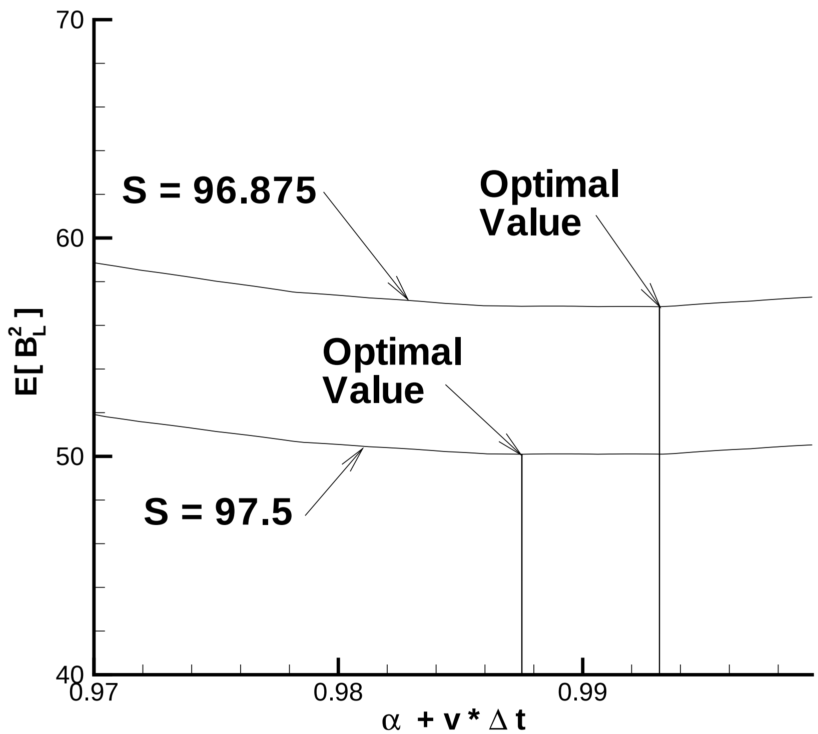

Value Surface $E_{v*}^{t=0}[\mathcal{B}_L^2]$ ($\mathcal{B} = -100$)

- Note $V_S \simeq V_\alpha \simeq V \simeq 0$ for large region. Similarity reduction $\rightarrow V_\mathcal{B} \simeq 0$

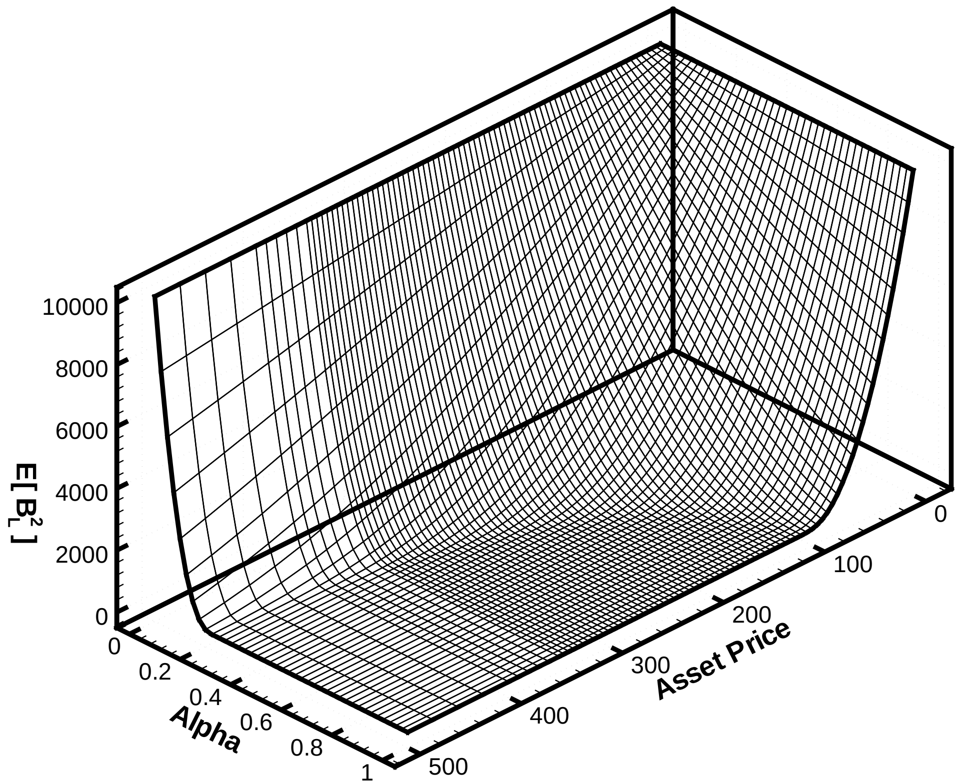

Optimal Strategy: Uniqueness

Recall HJB equation

$$V_\tau \;=\; \mathcal{L}V + \min_{v \in [v_{min}, v_{max}]} \Big[-vSf(v)V_{\mathcal{B}} + vV_\alpha + g(v)SV_S \Big] \;\; .$$

- If $V_S = V_{\mathcal{B}} = V_\alpha = 0$

$\hookrightarrow$ Optimal control can be any value $v \in [v_{min}, v_{max}]$.

Example: Non-uniqueness

Suppose we want a strategy which produces zero standard deviation
• Obvious method: sell immediately at infinite rate
$\hookrightarrow$ This gives zero standard deviation, but zero gain.

But this strategy is not unique. Another strategy:
• Do nothing until $T - \varepsilon$.
$\hookrightarrow$ Then sell at infinite rate.
• This strategy also produces zero standard deviation (and zero gain).

There are an infinite number of such strategies

Optimal Strategy: Discrete Trade Rates

Instead of allowing a continuous set of trade rates in $[v_{\min}, v_{\max}]$
$\hookrightarrow$ Allow only a set of twenty five possible trading rates in $[v_{\min}, v_{\max}]$

Solve HJB problem again.

Does this restriction to a fixed number of possible trading rates result in a large change in the efficient frontier?

Optimal Strategy: Discrete Trade Rates

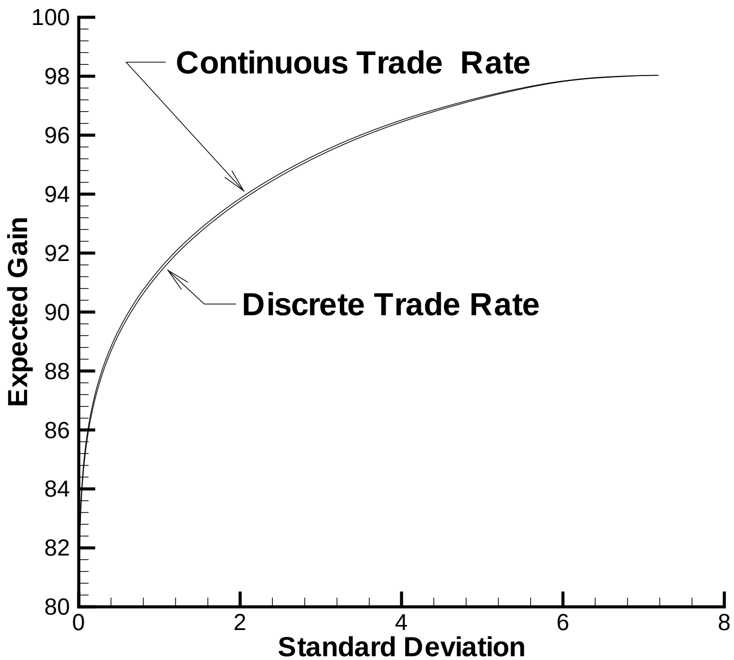

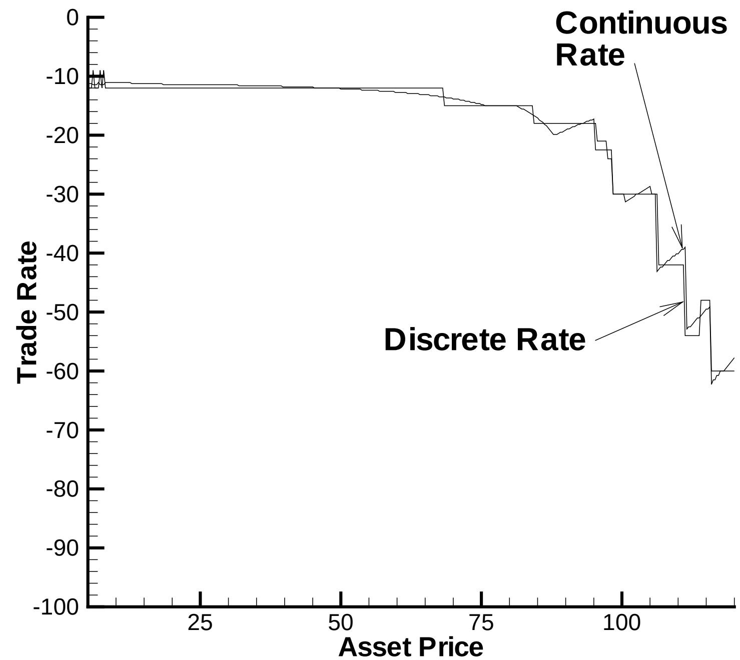

Clearly, many strategies which give almost the same efficient frontier!

Optimal Strategy: Discrete Trade Rates

The efficient frontier is very stable

- Appears to converge rapidly as mesh/timesteps refined.

The optimal trading rate is ill-posed.

- Strategy is non-unique in many cases.

This is actually useful in practice

- We get virtually the same efficient frontier for nearby strategies.
- Precise choice of trading strategy at start of trading not crucial.

Conclusions I

- One solve of the HJB equation, plus linear PDE, gives us the optimal (pre-committment) strategy for all points on the efficient frontier.
- Semi-Lagrangian solution of HJB PDE is independent (within reason) of

  ↪ Form of price impact functions

  ↪ Stochastic process for underlying asset (i.e. easy to add jumps, regime switching).
- We are finding the truly optimal mean-variance strategy

  ↪ For liquidation, if stochastic price moves higher, sell faster; if price drops, sell slower

Conclusions II

- Solution for efficient frontier is well-posed
  $\hookrightarrow$ Optimal strategy appears not well posed
  $\hookrightarrow$ Many strategies which give same frontier
  $\hookrightarrow$ Practical result: efficient frontier not sensitive to precise strategy!
- General selling rule:
  $\hookrightarrow$ Cut your gains, ride your losses.
- Question: is the pre-commitment mean-variance objective function what we want to do?

Disclaimer

This work was supported by a Morgan Stanley Equity Market Microstructure Research Grant. The views expressed herein are solely those of the authors, and not those of any other person or entity, including Morgan Stanley. Morgan Stanley is not responsible for any errors or omissions. Nothing in this presentation should be construed as a recommendation by Morgan Stanley to buy or sell any security of any kind.