ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Adaptive Arrival Price

Julian Lorenz (ETH Zurich, Switzerland)

Robert Almgren (Adjunct Professor, New York University)

Algorithmic Trading 2008, London
07. 04. 2008

Outline

- Evolution of algorithmic trading
- “Classic” arrival price algorithms
- Price adaptive algorithms
  - Single-update and multi-update strategies
  - Dynamic programming
  - Improved shortfall statistics and efficient frontiers
- Extensions
  - Bayesian adaptive trading with price trend

Electronic/Algorithmic Trading

- Use computers to execute orders
- Agency trading
  - Executing orders for clients
  - Investment decision is made
- Large and increasing fraction of total order flow
  - 92% of hedge funds and

    11% of all trades, 17% by 2007 (Tabb Group)
  - Expected 50% of traded volume by 2010 (The Economist, '07)

Evolution of Algorithmic Trading

Use computers to execute orders **(agency trading)**

Evolution

1. **VWAP**
   - Automatization of routine trading tasks
2. **Arrival Price**
   - Quantitative modeling and optimization
   - Market impact, volatility, alpha, *etc*
3. **Adaptivity**
   - Execution trajectory responds to market behavior in a variety of ways
   - **How to do this optimally?**

VWAP

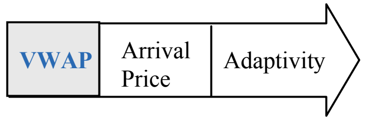

- Easy to understand and implement
  - “Spread trades out over time”
- Criticism
  - For large trades in illiquid securities, VWAP essentially reflects trade itself and provides little incentive for low-cost execution
  - Artificial: does not correspond to any investment goal

Arrival Price

Benchmark: Pre-trade (decision) price

If you could execute entire order instantly at this price, you would

Why trade slowly?

⇨ Reduce market impact

Why trade rapidly?

⇨ Minimize risk

⇨ Anticipated drift

- "Trader's Dilemma"
- Control statistical properties of shortfall
  - Risk-Reward tradeoff

Arrival Price: Efficient Strategies

Mean-Variance (Markowitz optimality): Strategies that minimize

- E for fixed V,
- V for fixed E, or
- E + λ V for risk aversion parameter λ

} Equivalent formulations!

❶ Minimal variance

❷ Minimal expected cost

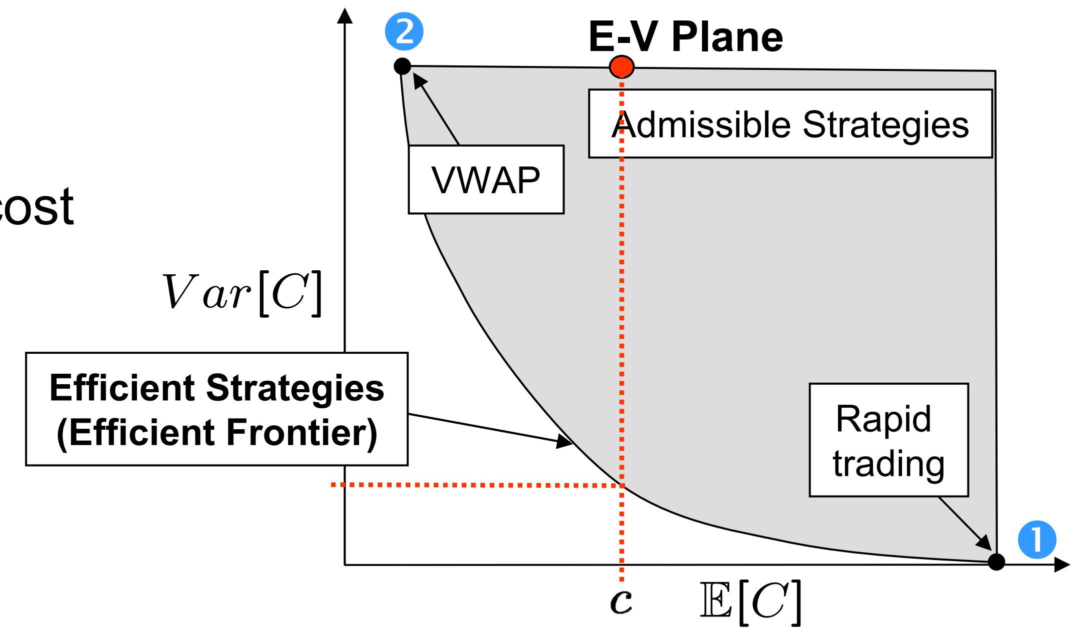

Arrival Price: Almgren/Chriss Trajectories

[Almgren, Chriss '00]: **Static trajectories** specified at **t=0**

- Efficient trajectories given by

$$x(t) = \frac{\sinh(\kappa(T-t))}{\sinh(\kappa T)} X$$

x(t) = stock holding at time t
T = time horizon
X = initial shares

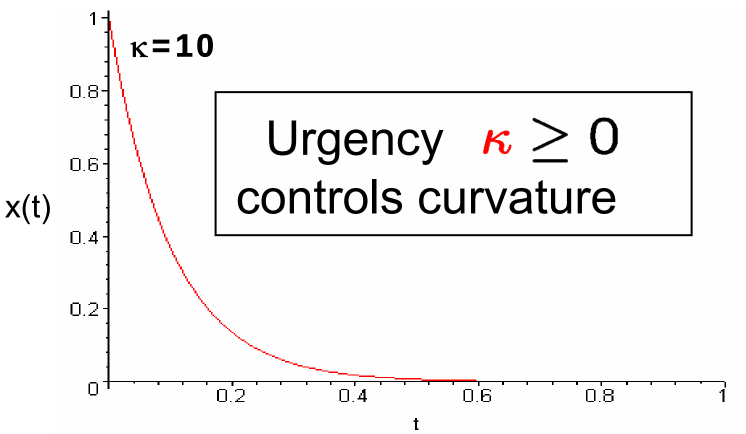

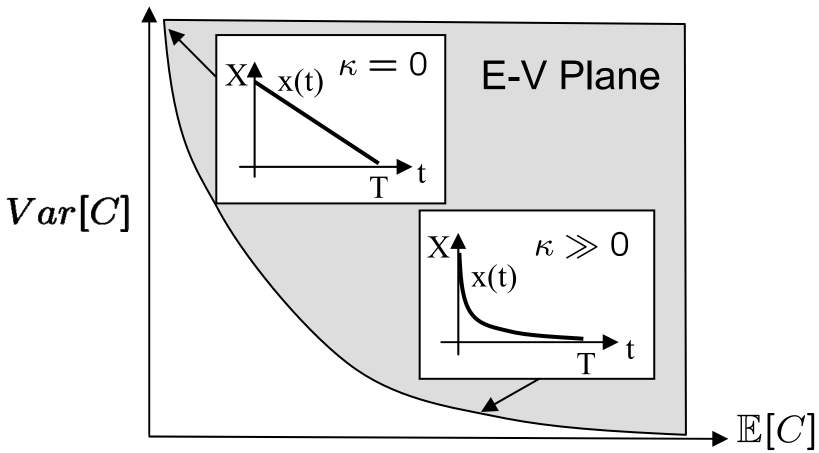

Adaptivity

Intuitively, adaptivity makes a lot of sense

1. Adapt to varying volume and volatility
   - trade more when liquidity is present
   - trade less when volatility risk is lower

2. Adapt to price movement: “scaling”
   - trade faster or slower when price move is favorable?

Scaling in response to price motions

Common wisdom: Depending on asset price process

- Mean-reversion $\Longrightarrow$ aggressive in the money **(AIM)**
- Momentum $\Longrightarrow$ passive in the money **(PIM)**
- Pure random walk $\Longrightarrow$ no response

[Almgren, L.]: AIM improves mean-variance tradeoff, especially

- for nonzero risk aversion (middle of frontier)
- and large portfolio sizes

1. **Single-Update:** [Algorithmic Trading III, Spring 2007]
   - Adapt strategy exactly once during trading
2. **Multi-Update:** [In preparation]
   - Any desired degree of precision

Intuition behind AIM for Arrival Price

- Introduce **intertemporal anti-correlation** between
  - investment gains/losses in first part of execution, and
  - trading costs (market impact) in second part of execution
- If make money in first part, spend parts on higher impact
- Higher impact = trade faster (i.e. reduces market risk)
- Trade this volatility reduction for expected cost reduction (mean-variance tradeoff!)
- Caveat: “Cap winners, let losers run” is deadly if real price process has momentum

1. Single-Update

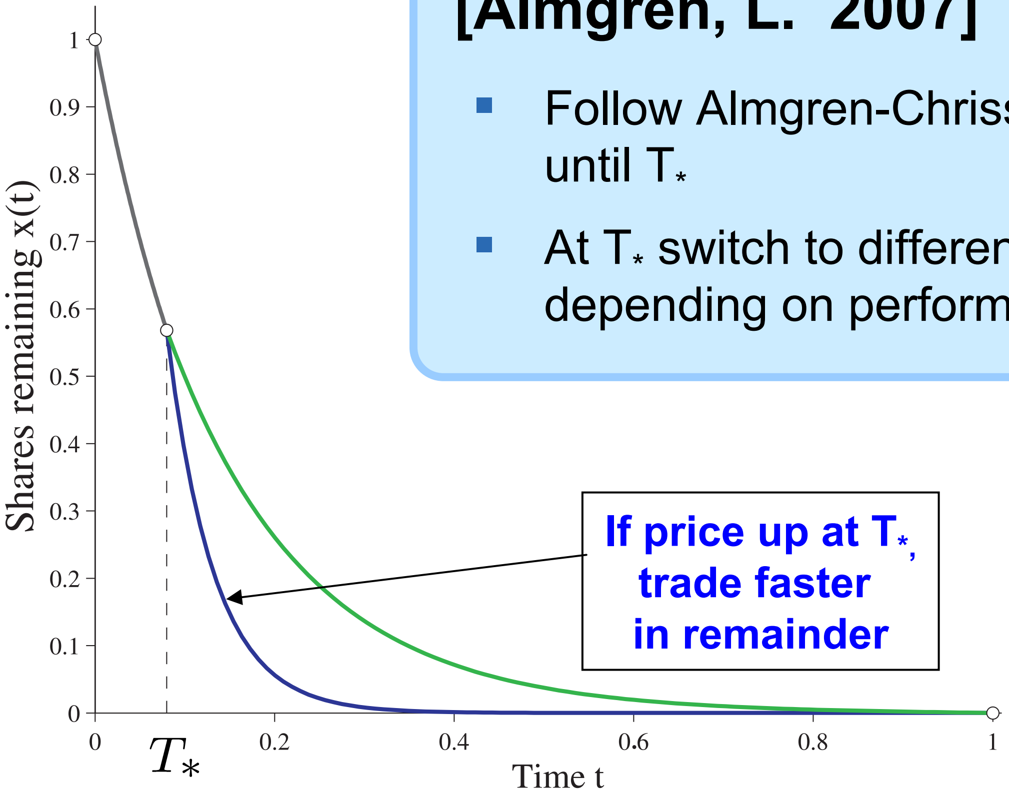

[Almgren, L. 2007]

- Follow Almgren-Chriss trajectory with urgency κ_0 until T_*
- At T_* switch to different urgency $\kappa_1,\dots,\kappa_n$ depending on performance up to T_*

Single-Update (cont.)

- Switch to κ_i if based on accumulated gain/loss at T_*
- $C_0, C_1, \ldots, C_n$ cost on first and second part

$$C_0 = \sigma \int_0^{T_*} x_0(t)\, dB(t) + \eta \int_0^{T_*} x_0'(t)^2\, dt \sim \mathcal{N}(E_0, V_0)$$

$$C_j = \sigma \int_{T_*}^{T} x_j(t)\, dB(t) + \eta \int_{T_*}^{T} x_j'(t)^2\, dt \sim \mathcal{N}(E_j, V_j)$$

- Explicit expressions for $E[C_i]$ and $Var[C_i]$
- Explicit expression for E and Var of composite strategy
  → Numerical optimization of criterion E +λ Var over $\kappa_0, \ldots, \kappa_n$

Single-Update: Numerical Results

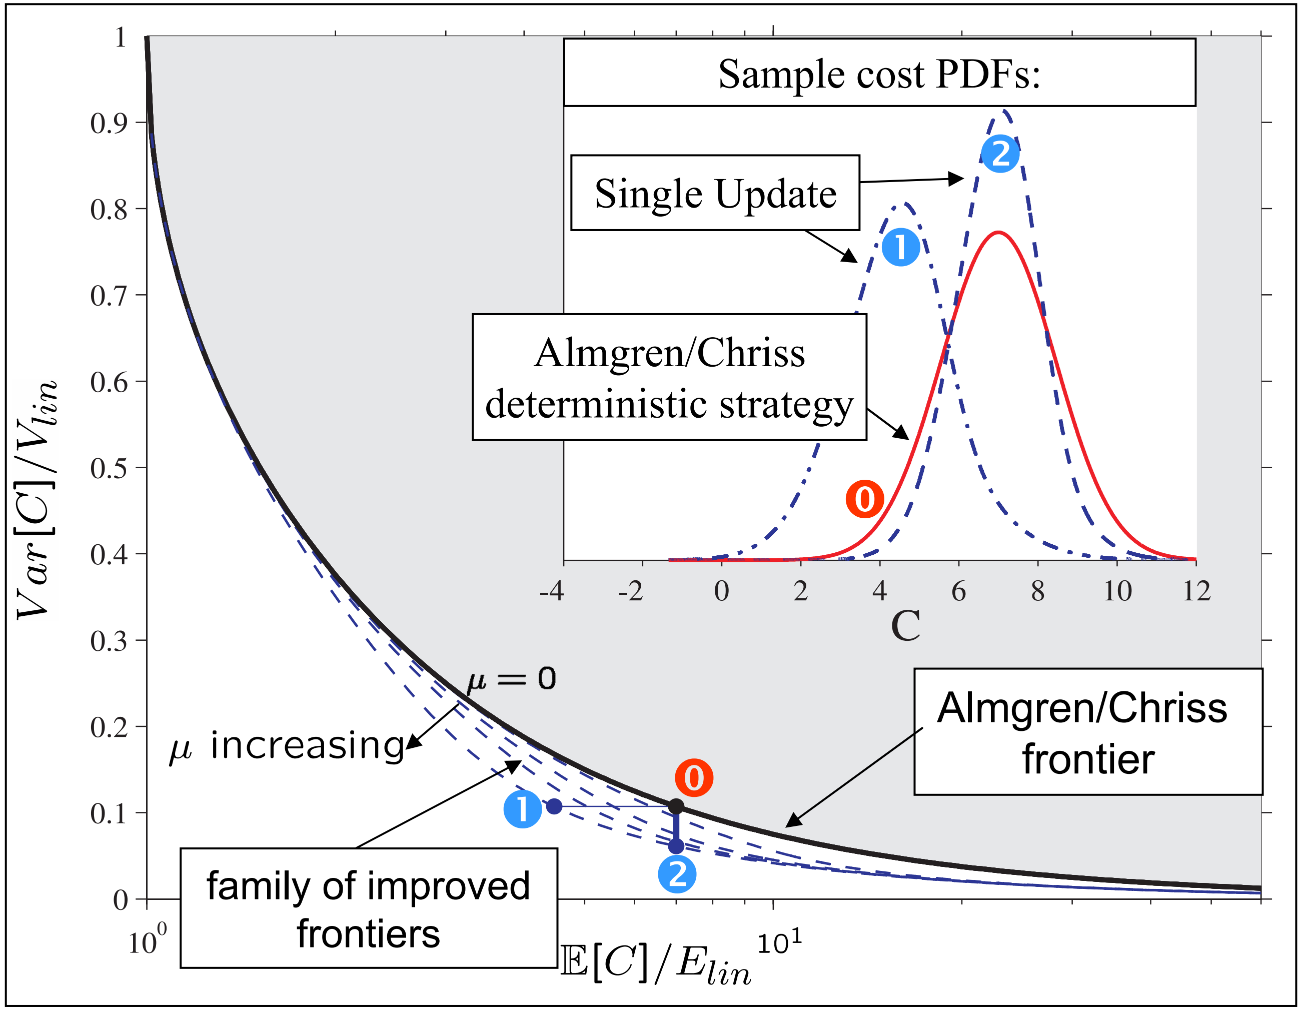

Family of frontiers parametrized by

market power μ

= measure of trade's ability to move the market compared to stock volatility

Larger relative improvement for large portfolios

$\mu \gg 0$ (i.e. $X \gg 0$)

AC static frontiers coincide with $\mu \to 0$

2. Multiple Updates: Dynamic Programming

- Single-update does NOT generalize to multiple updates
- "E + λ Var" not amenable to dynamic programming
  - Squared expectation in Var[X] = E[X²] - E[X]²

▶ **[Almgren, L. 2008]:** Define value function

$$J_k(x, c) = \text{"minimal variance for } k \text{ periods and } x \text{ shares s.t. expected cost at most } c\text{"}$$

$J_{k-1}(x, c)$ and optimal strategies for k-1 periods

Optimal Markovian one-step control

⇒

$J_k(x, c)$ and optimal strategies for k periods

▶ **Markov property** for mean-variance efficient strategies

Dynamic Programming (cont.)

We want to determine $J_k(x, c)$

Situation:
- **Sell program** (buy program analogously)
- k periods and x shares left
- limit for expected cost is c
- current stock price S
- next price innovation is $\sigma\xi \sim N(0,\sigma^2)$

Price impact function

Construct optimal strategy π for k periods

1. In **current period** sell y shares at $\tilde{S} = S - h(y)$
2. Use an efficient strategy π' for remaining $x - y$ shares in remaining **k-1 periods**

Dynamic Programming (cont.)

Note: y must be deterministic, but when we begin π', outcome of ξ is known, i.e. we may choose $\pi' = \pi'(\xi)$ **depending on** ξ

⇨ Specify $\pi'(\xi)$ by its expected cost z(ξ)

⇨ Strategy π defined by $y \in \mathbb{R}$ and control **function z(ξ)**

- Expressions for cost of strategy π conditional on ξ
  - $\mathbb{E}[C(\pi)\,|\,\xi]$ and $Var[C(\pi)\,|\,\xi]$
- Use the laws of total expectation and variance
  - $\mathbb{E}[C(\pi)] \;=\; \mathbb{E}[\,\mathbb{E}[C(\pi)\,|\,\xi]$
  - $Var[C(\pi)] \;=\; \mathbb{E}[\,Var[C(\pi)\,|\,\xi]] + Var[\,\mathbb{E}[C(\pi)\,|\,\xi]]$

Optimization of $(\mathbb{E}[C], Var[C])$ by means of y and $z(\xi)$

Dynamic Programming (cont.)

Value function recursion:

$$J_k(x,c) = \min_{(y,z)\in\mathcal{G}_k(x,c)} \Big\{ Var[z(\xi) - \sigma\xi(x-y)] + \mathbb{E}[J_{k-1}((x-y), z(\xi))] \Big\}$$

where

$$\mathcal{G}_k(x,c) = \left\{ (y,z) \;\middle|\; \begin{array}{l} 0 \le y \le x, \quad \mathbb{E}[z(\xi)] + \eta y^2 \le c \\ y \in \mathbb{R},\;\; z \in L^1(\Omega;\mathbb{R}) \end{array} \right\}$$

Control variable
new stock holding
(i.e. sell y in this period)

Control function
targeted cost as
function of next
price change ξ

(For linear price impact)

Solving the Dynamic Program

- Series of one-period optimization problems
- Each step: Determine optimal control
  - #shares to sell in next period
  - Target expected cost for remainder **as function z(ξ)** of next period stock price change
- No closed form solution → numerical approximation
- Nice property: **Convex constrained problems**

Behavior of Adaptive Strategy

Aggressive in the Money (AIM)

Formally: Control function z(ξ) is monotone increasing

Example: sell program

Spend windfall gains on increased impact costs to reduce total variance

Efficient Frontiers

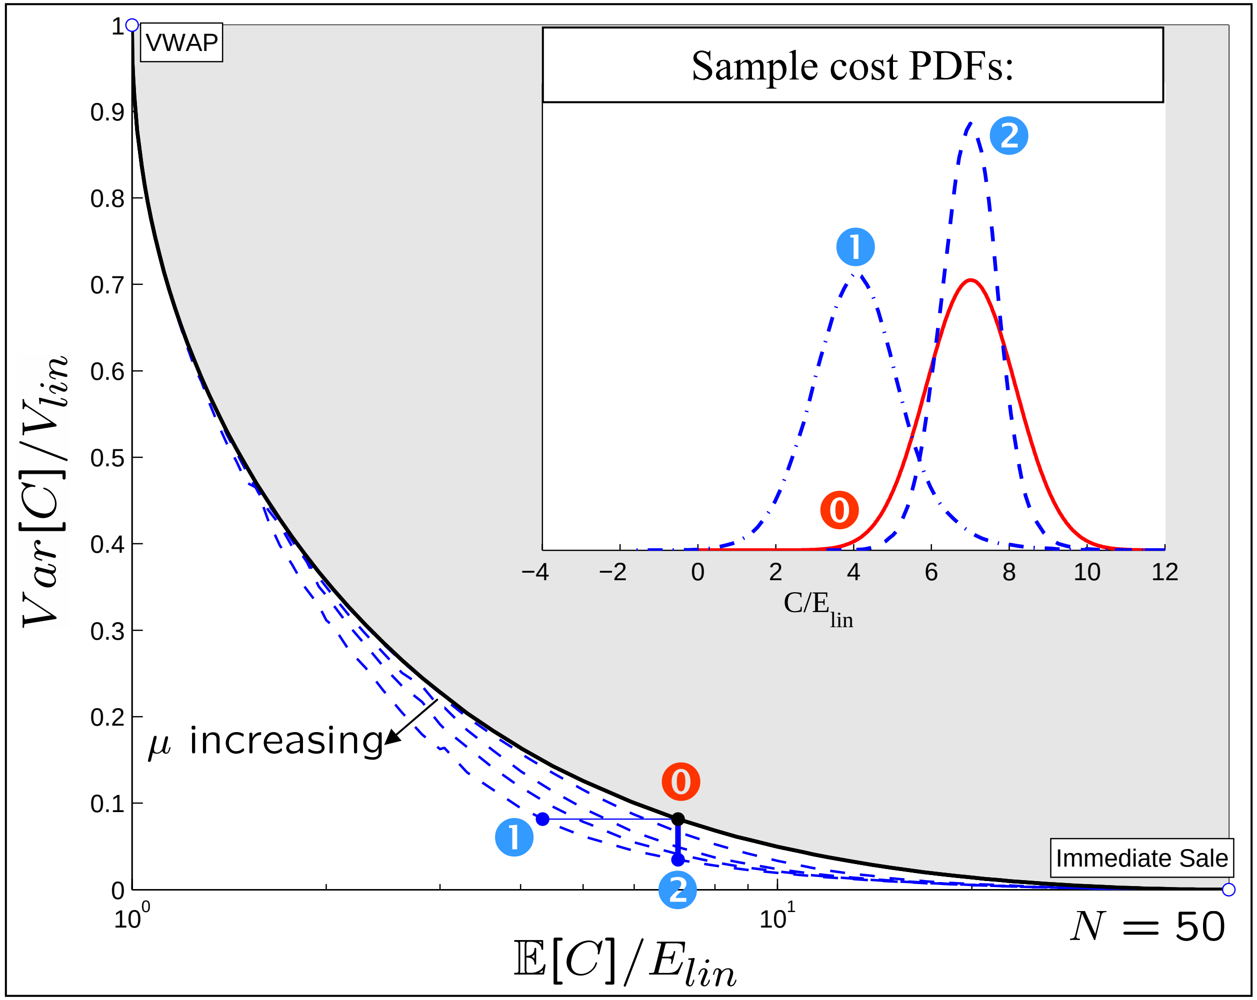

Similar results as for single-update

- **Larger relative improvement for large portfolios**
- **Market power μ**

Single Update vs. Multi-Update

- Full improvement by multi-update
- Significant improvement even by single-update
- Multi-update naturally more computational intensive
- Single-update offers good value for low computation

Other Criteria

- Instead of mean-variance tradeoff: Utility functions
  - Exponential utility (CARA): $u(y) = -\exp(-\alpha y)$
  - Power law utility: $u(y) = (y^{1-\gamma} - 1)/(1 - \gamma)$
  - etc.
- Optimal strategies are AIM or PIM depending on utility
  [Schied, Schöneborn '08]
- Advantage of mean-variance optimization
  - Clear and intuitive meaning
  - Corresponds to how trade results are reported in practice
  - Independent of client's wealth

Extensions

- Non-linear impact functions and cost models
- Multiple securities (program trading)
- Other asset price processes
  - Price momentum (drift)
  - Mean-Reversion
  - Non-Gaussian returns
- etc.

Trading with Price Trend

"Daily Trading Cycle": institutional traders make decisions overnight and trade across the following day

- Price momentum, if large net positions being traded

Market Model

- Stock price: Brownian motion with *unknown* drift α

$$S(t) = S_0 + \alpha t + \sigma B(t) \quad \text{with} \quad \alpha \sim \mathcal{N}\left(\bar{\alpha}, \nu^2\right)$$

- Prior estimate for drift, update this belief using price observations
- Temporary and permanent market impact

[Almgren, L. '06]: "Bayesian Adaptive Trading", Journal of Trading

Optimal risk-neutral strategy (buy program)

x(t) = trading trajectory (shares remaining) in continuous time
v(t) = instantaneous trading rate

[Almgren, L. 2007]: Optimal dynamic strategy given by instantaneous trade rate

$$v_*(t) = \frac{x(t)}{T-t} + \hat{\alpha}(t, S(t)) \cdot \frac{T-t}{4\eta}$$

T = time horizon
η = linear market impact

- $\hat{\alpha}(t, S(t))$: Best estimate of drift using prior and S(t)
- Trade rate of re-computed static trajectory with current best drift estimate
  - Locally optimal myopic strategy = Global optimal solution
  - Highly dynamic

Trading with Price Trend: Examples

Buy-program with significant upwards drift prior

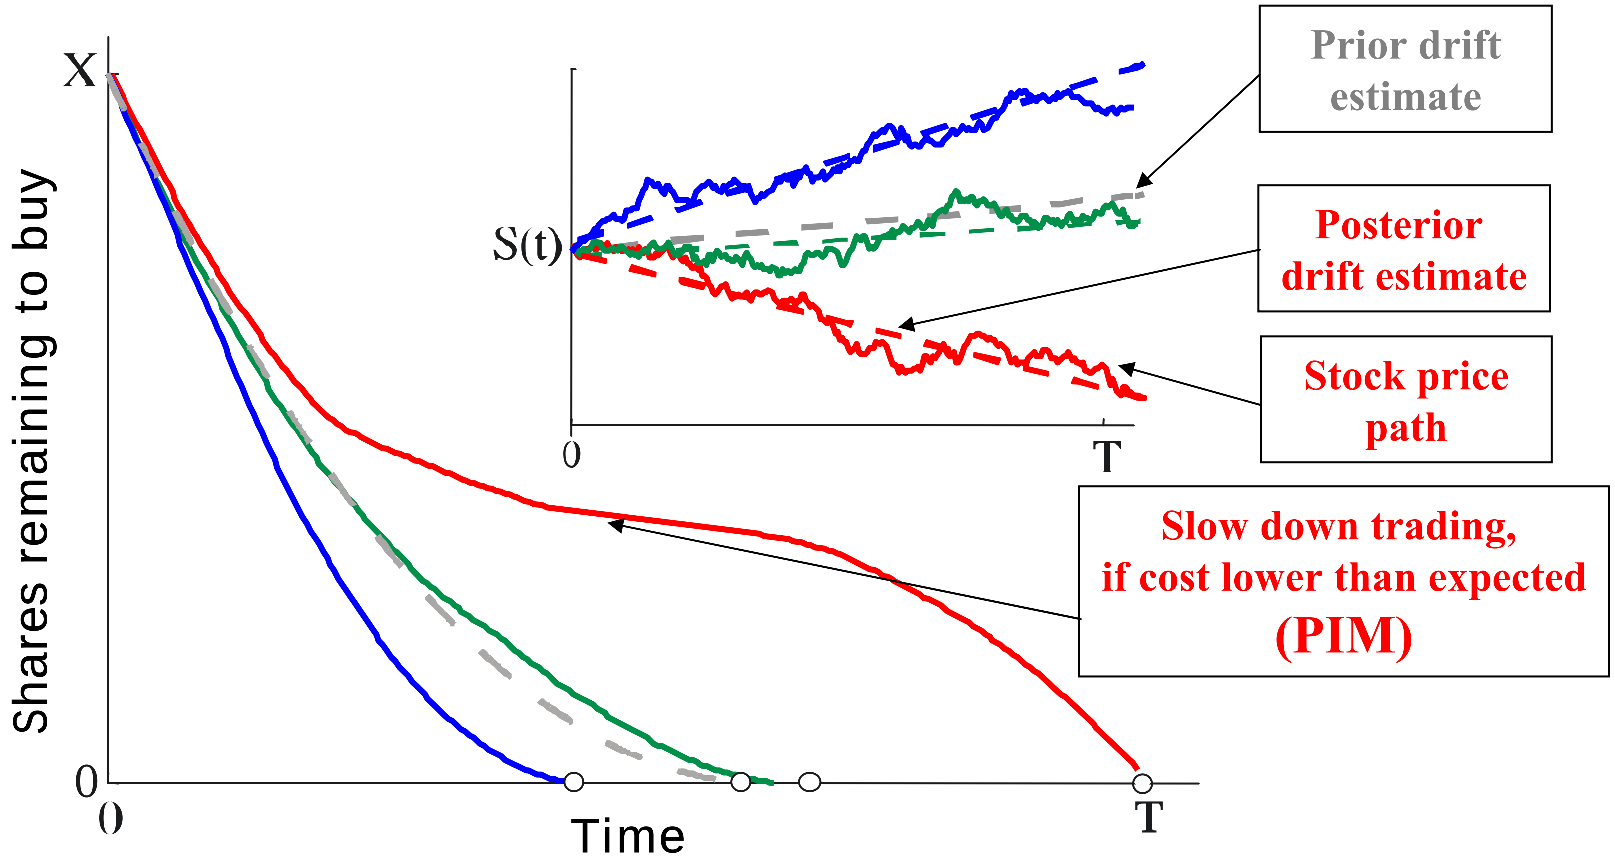

Conclusion

- Adaptivity is the new frontier of algorithmic trading
- Single-update and multi-update arrival price algorithms
  - Pure random walk, no momentum or reversion
  - Straightforward mean-variance criterion
  - Dynamic programming approach
  - Strategies are “aggressive-in-the-money“
- Substantially *better than static Almgren/Chriss trajectories*
  - Improved efficient frontiers
  - Relative improvement bigger for large portfolios
  - New market power parameter μ

Thank you very much for your attention!

Questions?