Algorithmic Challenges in Modern Financial Markets

Michael Kearns
Computer and Information Science
University of Pennsylvania
and
Equity Strategies Group
Lehman Brothers New York

Acknowledgements: Eyal Even-Dar, Elliot Feng, Sham Kakade, Yishay Mansour, Yuriy Nevmyvaka, Luis Ortiz, Cenk Ural, Jenn Wortman

ACM EC Tutorial
Ann Arbor, MI
June 12, 2006

A Stock Market Field Guide
(The “Biodiversity” of Wall Street)

- Retail traders
  - individual consumers
- “Buy” side
  - large institutional traders: portfolio managers; mutual and pension funds; endowments
  - often have precise metrics and constraints; e.g. tracking indices
  - percentage-based management fee
- “Sell” side
  - brokerages providing trading/advising/execution services
  - “program trading” → “algorithmic trading”: automated strategies for optimized execution
  - profit from commissions/fees
- Market-makers and specialists
  - risk-neutral providers of liquidity
  - highly regulated
  - profit from the “bid-ask bounce”; averse to strong directional movement
  - automated market-making strategies in electronic markets
- Hedge funds and proprietary trading
  - groups attempting to yield “outsized” returns on private capital (= beat the market)
  - can take short positions
  - highly unregulated; starting to see institutional investment
  - heavy quant consumers: “statistical arbitrage”, modeling, algorithms
  - typically take management fee and 20% of profits
- All have different goals, constraints, time horizons, technology, data, connectivity…

Where are the Algorithmic Challenges?

- Need:
  - precisely specified constraints (inputs and outputs)
  - measures of performance
  - data
- Two important areas:
  - Part I: Market Microstructure and Optimized Execution
  - Part II: Proprietary Trading and (Generalized) Portfolio Optimization

Part I:
Market Microstructure and Optimized Execution

Questions of Enduring Interest

- How do (stock) prices “evolve”? How can we model this evolution?
  - classical random walk, diffusion models + drift
    - many recent empirical challenges [Lo & MacKinlay; Brock et al.]
  - autoregressive time series models
    - AR1, ARCH, GARCH, etc. → generalized Ito model
  - computer science: adversarial/worst-case price sequences
    - algorithms analyzed w.r.t. competitive ratios, regret
- Can we design “adaptive” or “learning” algorithms for:
  - executing difficult/large trades?
  - predicting and profiting from movements of prices?
- Models generally ignore market mechanism and liquidity issues
  - at least in part because the data was unavailable and unreliable
- This is changing rapidly… and presents challenges & opportunities

Part I Outline

- Market Microstructure and Optimized Execution
- Competitive Analysis for VWAP and Limit Order Trading
- Reinforcement Learning for Optimized Execution
- (In)Stability Properties of Limit Order Dynamics

Background on Market Microstructure and Optimized Trade Execution

Background on Market Microstructure

- Consider a typical exchange for some specific security
- Limit order: specify price (away from the market)
- (Partially) Executable orders are filled immediately
  - prices determined by standing orders in the book
  - one order may execute at multiple prices
- Non-executable orders are placed in the buy or sell book
  - sorted by price; top prices are the bid and ask
- Market order: limit order with an extreme price
- Full order books now visible in real time
- What are they good for?

refresh | island home | disclaimer | help

MSFT

GET STOCK

MSFT go

Symbol Search

LAST MATCH		TODAY'S ACTIVITY	
Price	24.0700	Orders	52,983
Time	14:57:07.72	Volume	10,243,212

BUY ORDERS		SELL ORDERS	
SHARES	PRICE	SHARES	PRICE
500	24.0620	500	24.0690
6,000	24.0610	500	24.0690
5,000	24.0600	500	24.0700
100	24.0600	200	24.0800
1,100	24.0550	1,981	24.0900
100	24.0500	412	24.0900
5,000	24.0500	3,000	24.0980
200	24.0500	500	24.1000
3,294	24.0500	100	24.1200
1,000	24.0500	2,800	24.1400
3,000	24.0430	5,000	24.1400
100	24.0400	1,000	24.1400
5,503	24.0400	5,000	24.1500
2,100	24.0300	400	24.1600
2,800	24.0300	1,000	24.1700
(412 more)		(694 more)	

As of 14:57:16.178

Optimized Trade Execution

- Canonical execution problem: sell V shares in T time steps
  - must place market order for any unexecuted shares at time T
  - also known as “one-way trading” (OWT)
  - trade-off between price, time… and liquidity
- Problem is ubiquitous
- Multiple performance criteria:
  - Maximum Price:
    - compare revenue to max execution price
    - O(log(R)) competitive ratios in infinite liquidity, adversarial price model
      - R = a priori bound on ratio of max to min execution price
      - [El-Yaniv, Fiat, Karp & Turpin]
  - Volume Weighted Average Price (VWAP):
    - compare to per-share average price of executions
    - widely used on Wall Street; reduces risk sources to execution
    - by definition, must track prices and volumes
  - Implementation Shortfall:
    - compare per-share price to mid-spread price at start of trading interval
    - an unrealizable ideal

Algorithms for VWAP and Limit Order Trading

[Kakade, K., Mansour, Ortiz ACM EC 2004]

An Online Microstructure Model

- Market places a sequence of price-volume limit orders:
  - M = (p_1,v_1),(p_2,v_2),…,(p_T,v_T)   (+ order types)
  - possibly adversarial; also consider various restrictions
- Algorithm is allowed to interleave its own limit orders:
  - A = (q_1,w_1),(q_2,w_2),…,(q_T,w_T) (+ order types)
- Merged sequence determines executions and order books:
  - merge(M,A) = (p_1,v_1), (q_1,w_1),…, (p_T,v_T), (q_T,w_T)
  - now have complex, high-dimensional state
    - how to simplify?

refresh | island home | disclaimer | help

MSFT

GET STOCK

MSFT go

Symbol Search

LAST MATCH		TODAY'S ACTIVITY	
Price	24.0700	Orders	52,983
Time	14:57:07.72	Volume	10,243,212

BUY ORDERS		SELL ORDERS	
SHARES	PRICE	SHARES	PRICE
500	24.0620	500	24.0690
6,000	24.0610	500	24.0690
5,000	24.0600	500	24.0700
100	24.0600	200	24.0800
1,100	24.0550	1,981	24.0900
100	24.0500	412	24.0900
5,000	24.0500	3,000	24.0980
200	24.0500	500	24.1000
3,294	24.0500	100	24.1200
1,000	24.0500	2,800	24.1400
3,000	24.0430	5,000	24.1400
100	24.0400	1,000	24.1400
5,503	24.0400	5,000	24.1500
2,100	24.0300	400	24.1600
2,800	24.0300	1,000	24.1700
(412 more)		(694 more)	

As of 14:57:16.178

What Can Be Done?

- Maximum Price:
  - O(log(R)) inf. liquidity model → O(log(R)log(V)) in microstructure model
  - quantifies worst-case market impact of large trades
  - if p_1 > p_2 >… are execution prices, randomly "guess" max{kp_k}
  - note: optimal offline algorithm unknown!
- VWAP:
  - O(log(Q)) in microstructure
    - Q = ratio of max to min total executed volume on sequence
    - Q often small empirically; can exploit (entropic) distributional features
  - Better: trade V shares over γV executed shares, **γ > 1**
    - VWAP "with volume" instead of "with time"
  - Can approach competitive ratio of 1 for large V !
  - Sketch of algorithm/analysis:
    - divide time into equal (executed) volume intervals I_1, I_2,…
    - place sell order for 1 share at ~ (1-ε)^k nearest VWAP_j
    - if all orders executed, are within (1-ε) of overall VWAP
    - can't "strand" more than one order at any given price level
    - optimize ε
- None of these algorithms "look" in the order books!

Reinforcement Learning for Optimized Trade Execution

[Nevmyvaka, Feng, K. ICML 2006]

RL for Optimized Execution

- Basic idea: execution as state-based stochastic optimal control
  - state: time and shares remaining… what else?
  - actions: position(s) of orders within the book
  - rewards: prices received for executions
  - stochastic: because same state may evolve differently in time
- This work: large-scale application of RL to microstructure
- Related work:
  - Bertsimas and Lo
  - Coggins, Blazejewski, Aitken

“No Impact” State Factorization

Full OB State:

OB execution simulation → reward (share prices)

OB State Features: OB(T) → OB(T-1) → OB(T-2) →

What OB features?

Massive saving of data and computation… Will it work?

Policy: a(T) a(T-1) a(T-2)

Action: limit price for remaining volume

Private State: (T,V) → (T-1,V’) → (T-2,V’’) →

Training only, do full OB sim on test data

Experimental Details

- Stocks: AMZN, NVDA, QCOM (varying liquidities)
- V = 5K and 10K shares
  - divided into 1, 4 or 8 levels of observed discretization
- T = 2 and 8 mins
  - divided into 4 or 8 decision points
- Explored a variety of OB state features
- Learned optimal strategy on 1 year of INET training data
- Tested strategy on subsequent 6 months of test data
- Objective function:
  - basis points compared to all shares at initial spread midpoint
    - implementation shortfall; an unattainable ideal (infinite liquidity assumption)
- Same basic RL framework can be applied much more broadly
  - e.g. “market-making” strategies [Chan, Kim, Shelton, Poggio]

A Baseline Strategy: Optimized Submit-and-Leave

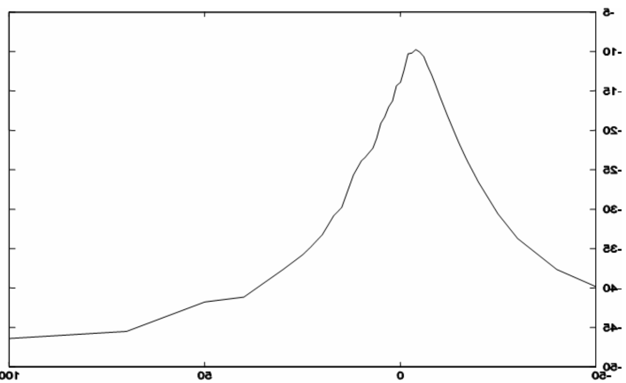

Trading Cost vs. Limit Price

deep in OB

M.O. at start

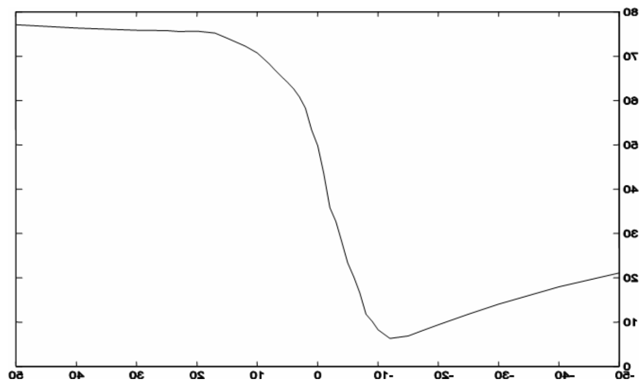

Risk vs. Limit Price

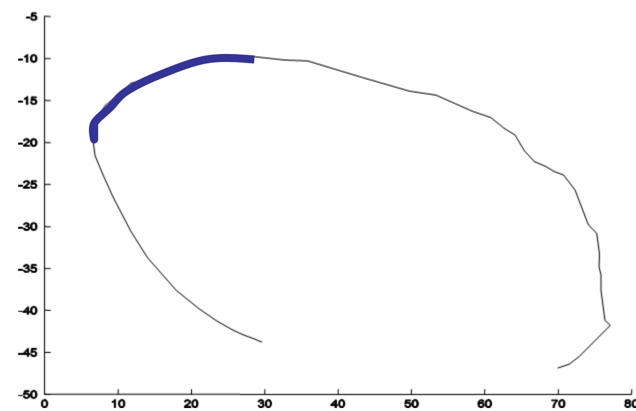

Efficient Frontier

[Nevmyvaka, K., Papandreou, Sycara IEEE CEC 2005]

Experimental Results

Private State Variables Only: Time and Inventory Remaining

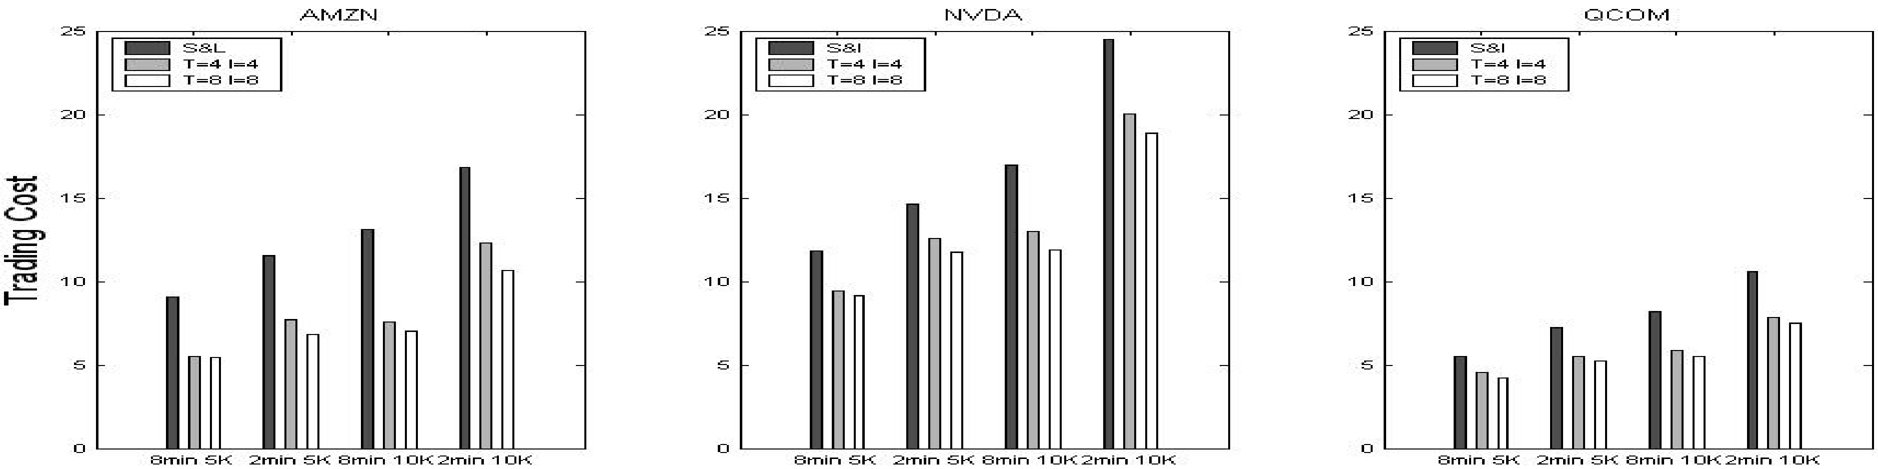

Average Improvement Over Optimized Submit-and-Leave

T=4 I=1	27.16%	T=8 I=1	31.15%
T=4 I=4	30.99%	T=8 I=4	34.90%
T=4 I=8	31.59%	T=8 I=8	35.50%

Strategy Visualization (10K, 2min)

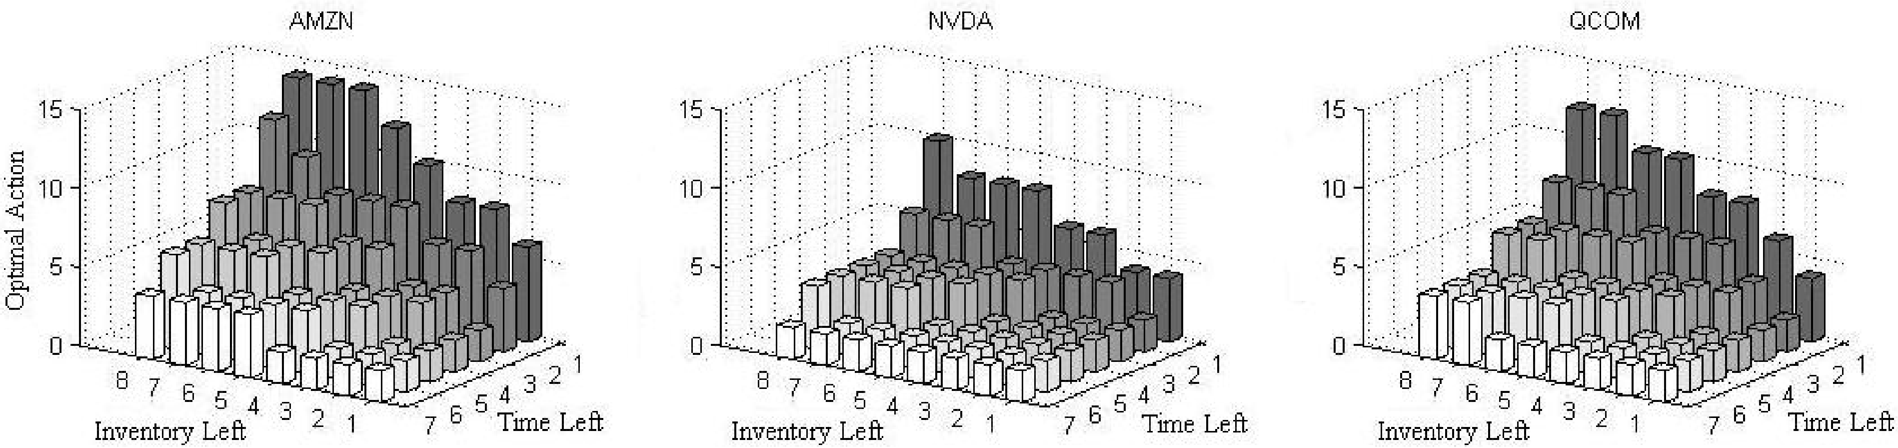

General shape is intuitive, but (stock-specific) numerical optimization matters!

Q-Values: Trading Costs vs. Actions (10K, 2min)

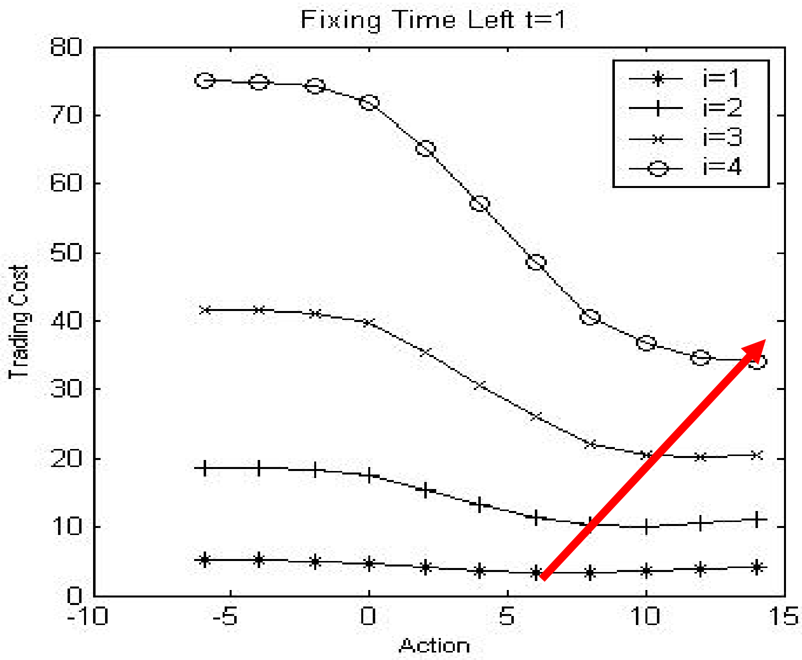

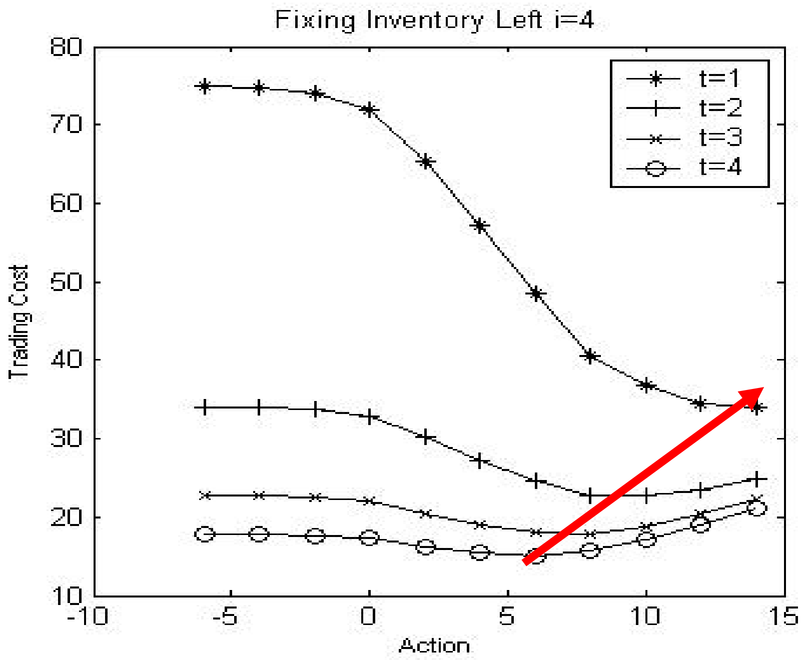

Improvement From Order Book Features

Bid Volume	-0.06%	Ask Volume	-0.28%
Bid-Ask Volume Misbalance	0.13%	Bid-Ask Spread	**7.97%**
Price Level	0.26%	Immediate Market Order Cost	**4.26%**
Signed Transaction Volume	**2.81%**	Price Volatility	-0.55%
Spread Volatility	1.89%	Signed Incoming Volume	0.59%
Spread + Immediate Cost	**8.69%**	Spread+ImmCost+Signed Vol	**12.85%**

Strategy Visualization II

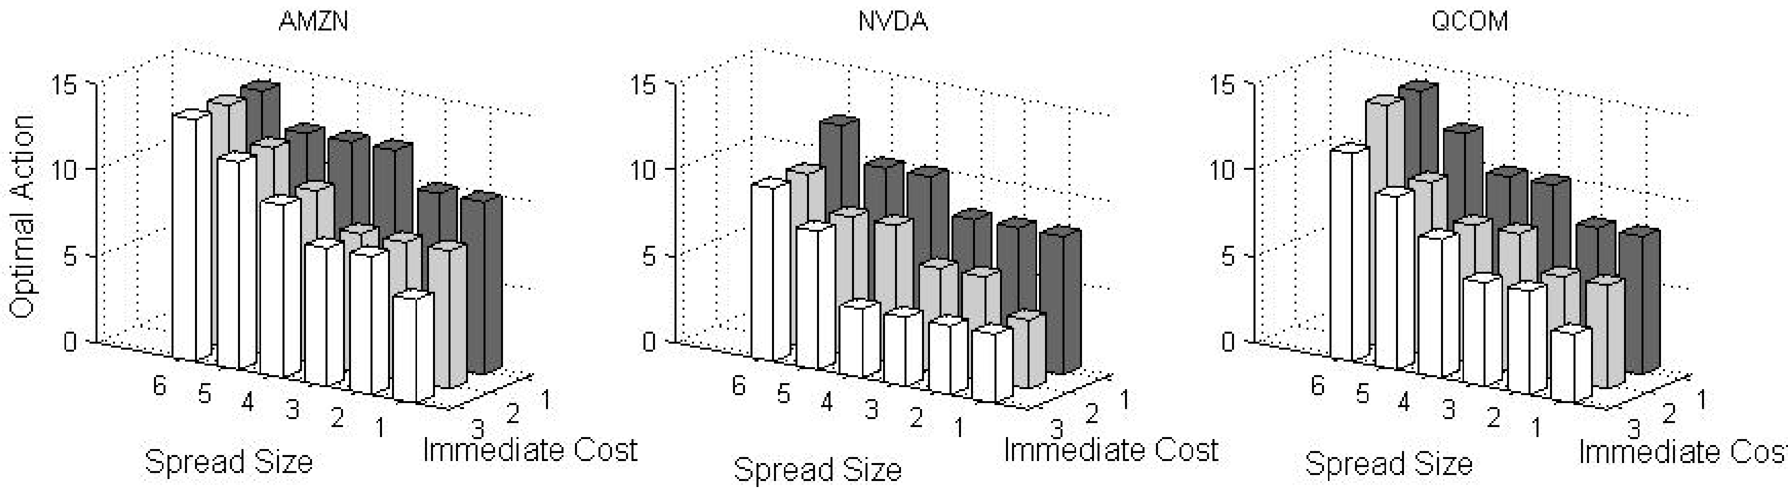

Q-Values: Trading Costs vs. Actions

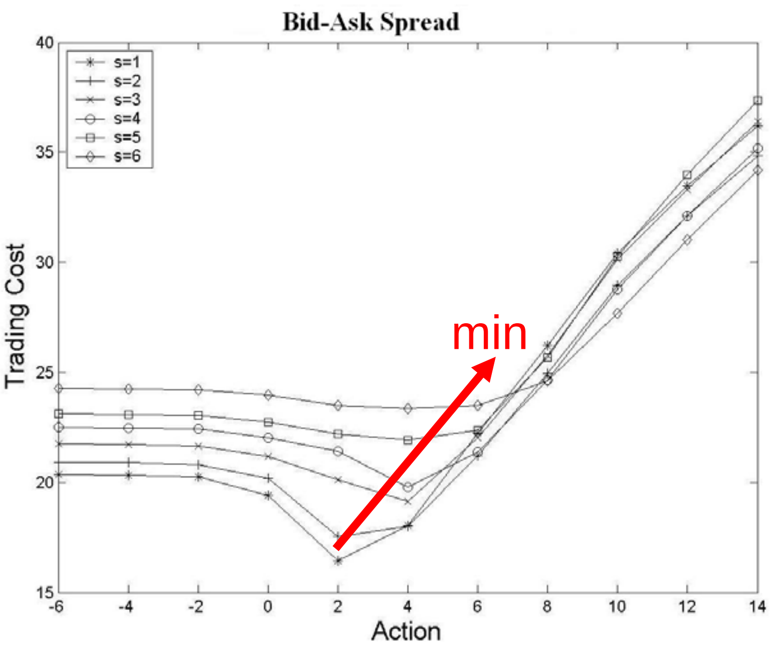

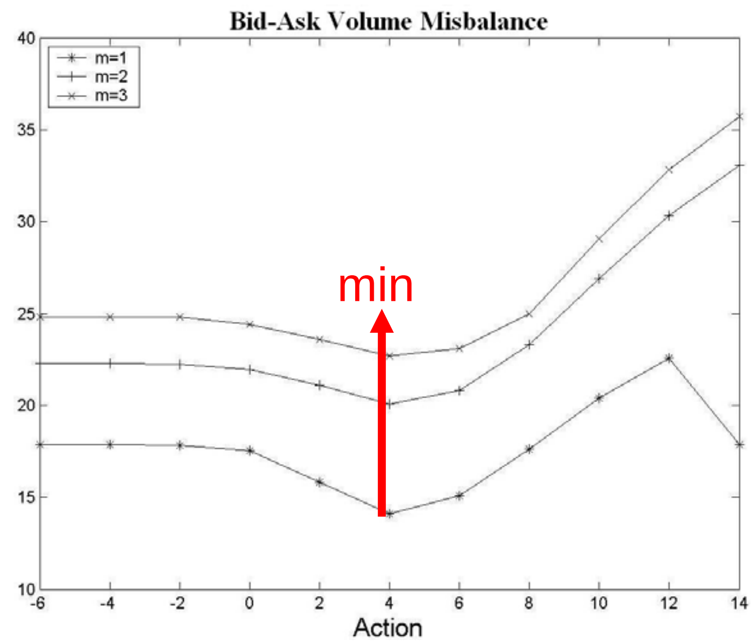

predictive and actionable

predictive but not actionable

(In)Stability Properties of Limit Order Dynamics

[Even-Dar, Kakade, K., Mansour ACM EC 2006]

“Backtesting” of Trading Strategies

- Theory and experiments describe so far:
  - assume access to limit order data (historical or “live”)
  - reconstruct complete order books at each point in time
  - insert hypothetical limit orders into the stream
    - competitive analysis: sequence of “market” limit orders arbitrary but fixed in advance
    - RL experiments: limit order data was historical
  - simulate forward the execution of the hypothetical orders
- Faithfully simulate the mechanical aspects of market impact
- What about the reactive or “psychological” aspects?
- Formalize as a question about dynamical stability:
  - Make various assumptions about how future orders do or do not react to the past
  - *Can tiny perturbations of the limit order sequence cause dramatic future change?*
  - Butterfly Effects and Chaos
- Of basic interest to any backtesting process… and thus to ML in finance

Two Models of Market Impact

- Both models deal with arbitrary, fixed sequences… but of what?
- Absolute model:
  - the model assumed so far
  - market given by a sequence of “absolute” limit order prices (one share each)
  - e.g. M = (p_1,buy),(p_2,buy),(p_3,sell),…
  - order books constructed from sequence M
  - “mechanical” impact only
  - motivation:
    - traders with “inherent” valuations
    - traders with slow time scales, long investment horizons, poor microstructure access
- Relative model:
  - market given by a sequence of limit order prices relative to current bid & ask
  - e.g. M’ = (d_1,buy),(d_2,buy),(d_3,sell),…
  - construct order books & actual prices in concert with each other
    - e.g. limit price p_2 = current bid + d_2; limit price p_3 = current ask + d_3; etc.
  - crude form of “psychological” or “reactive” impact
  - motivation:
    - traders “looking for a bargain”; trading off time for price
    - “penny-jumping”, optimized execution
- *How do these models differ?*

Stability

- Consider sequences in the two models:
  - absolute: M = (p_1,type_1),(p_2,type_2),…
  - relative: M' = (d_1,type_1),(d_2,type_2),…
- Now consider a small, arbitrary modification to each
  - e.g. deleting or adding a single order
    - (p_i,type_i) from M, (d_i,type_i) from M'
    - think of this as "our" action
- How much can such a change alter basic properties of the sequence?
  - stability = small change not amplified with time
  - instability = small change greatly amplified
- Absolute model: Every "reasonable" property stable!
  - volume executed, VWAP, closing price,…
  - note: must still be careful; some bounds depend on spread of M
  - generalizes to larger modifications, other types
- Relative model: Most properties highly unstable!
  - can find sequences (with bounded spread) such that single deletion causes arbitrarily large changes in volume executed, VWAP, closing price,…

Absolute Model Stability

- <B,S> = original buy and sell books (at some point in simulation)
- <B’,S’> = modified buy and sell books (at the same point)
- Introduce “meta-states” with small “edit distance” between simulations
- E.g. meta-state where B = B’ and S U {s’} = S’ U {s} for some s <> s’
- Main technical lemma establishes:

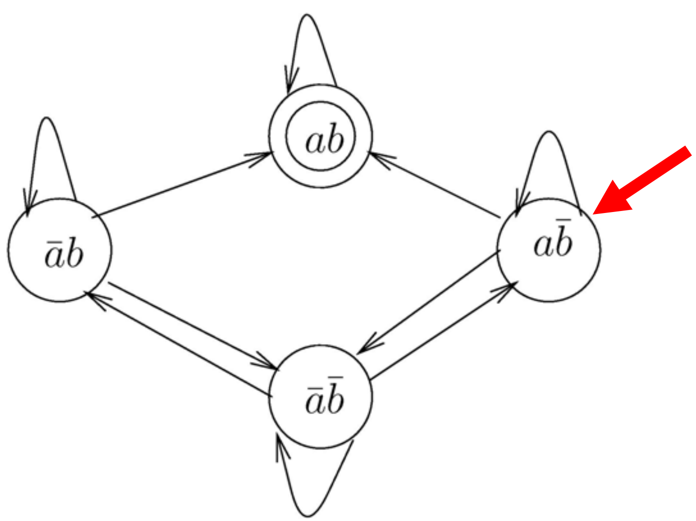

Figure 1: Diagram representing the set S of stable states and the possible movements transitions in it after the change.

Simulations

0.0025%

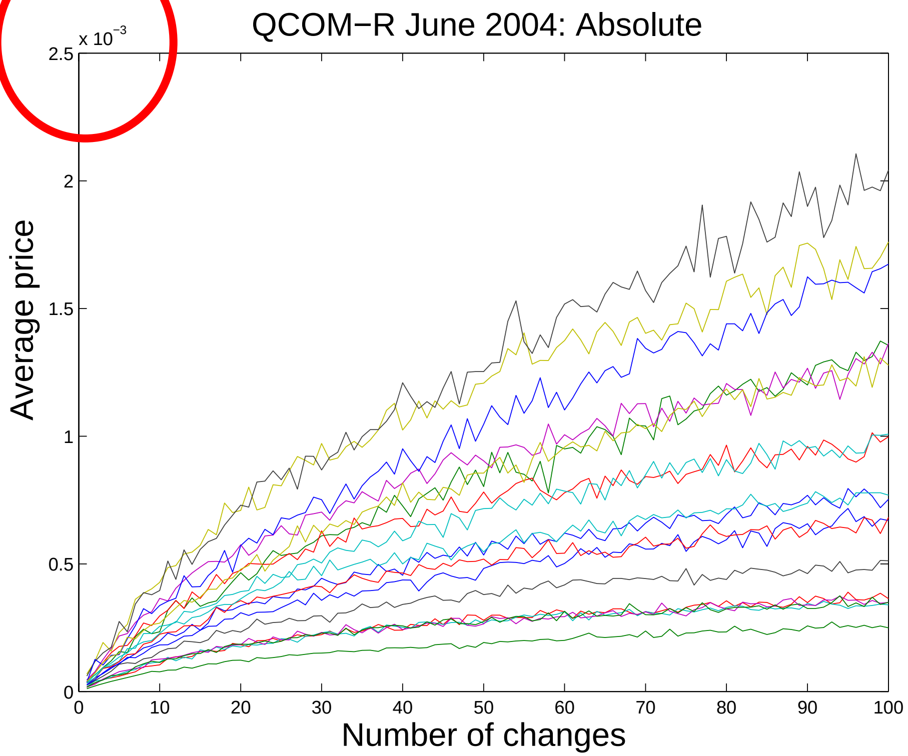

0.7%

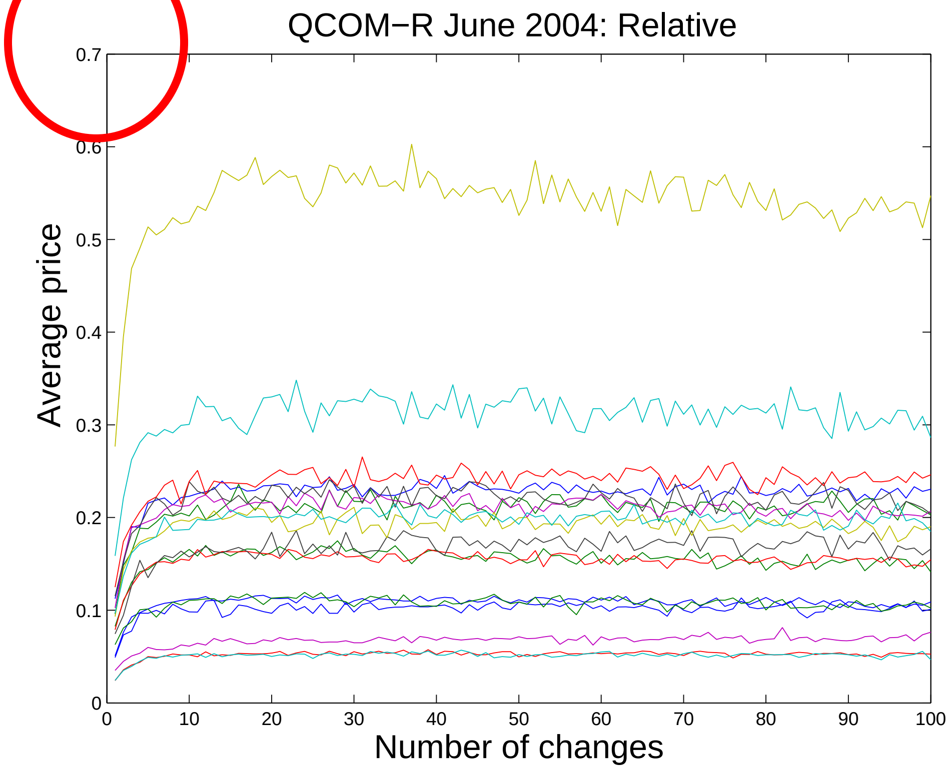

% change in VWAP vs. #changes: Absolute model

% change in VWAP vs. #changes: Relative model

A Mixture Model

fraction α of absolute traders, 1-α of relative traders

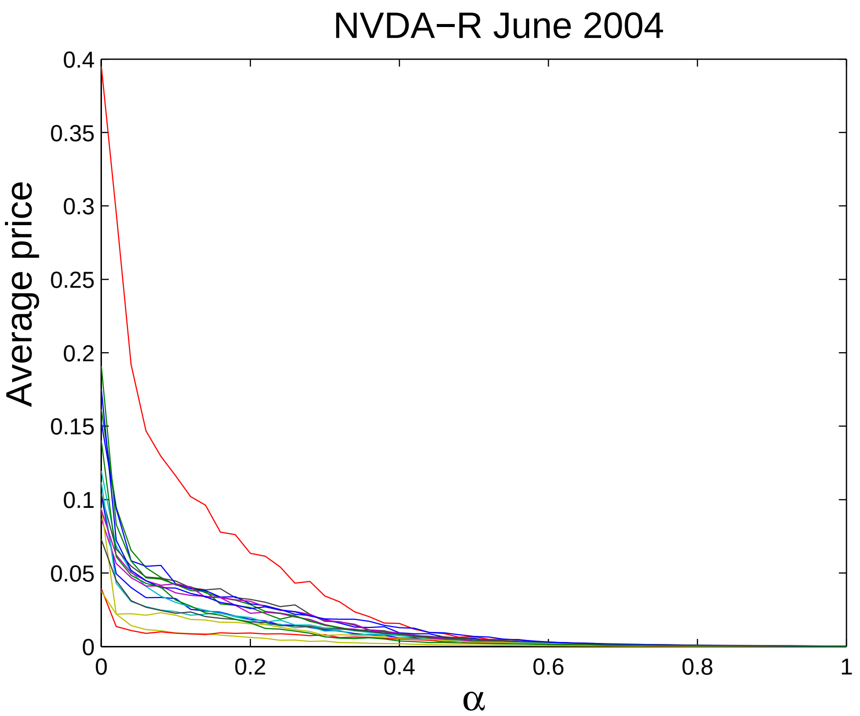

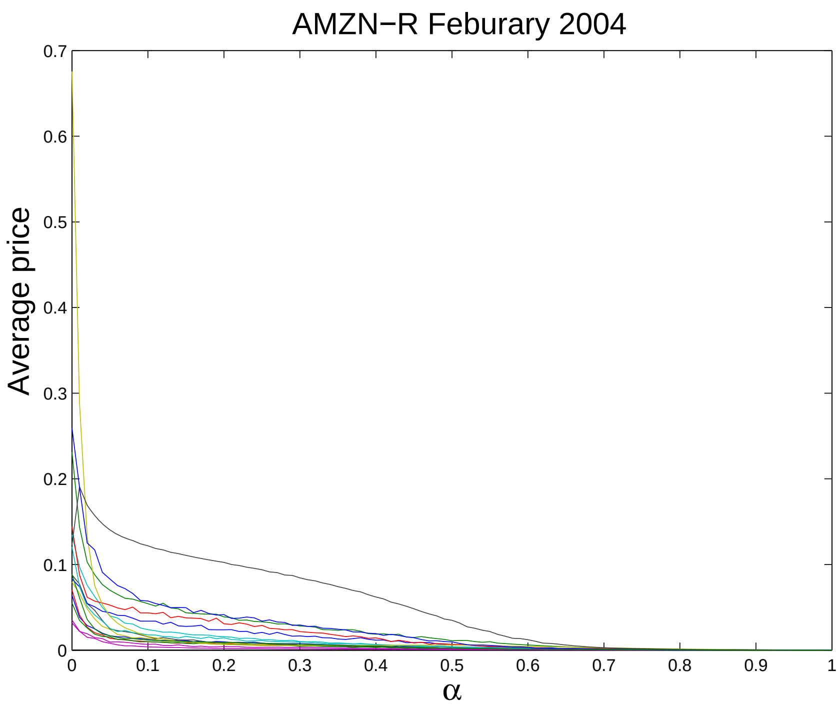

% change in VWAP vs. α,
single order deletion

Part II:
Proprietary Trading and
(Generalized) Portfolio Optimization

A Stock Market Field Guide
(The "Biodiversity" of Wall Street)

- Retail traders
  - individual consumers
- "Buy" side
  - large institutional traders: portfolio managers; mutual and pension funds; endowments
  - often have precise metrics and constraints; e.g. tracking indices
  - percentage-based management fee
- "Sell" side
  - brokerages providing trading/advising/execution services
  - "program trading" → "algorithmic trading": automated strategies for optimized execution
  - profit from commissions/fees
- Market-makers and specialists
  - risk-neutral providers of liquidity
  - highly regulated
  - profit from the "bid-ask bounce"; averse to strong directional movement
  - automated market-making strategies in electronic markets
- Hedge funds and proprietary trading
  - groups attempting to yield "outsized" returns on private capital (= beat the market)
  - can take short positions
  - highly unregulated; starting to see institutional investment
  - heavy quant consumers: "statistical arbitrage", modeling, algorithms
  - typically take management fee and 20% of profits
- All have different goals, constraints, time horizons, technology, data, connectivity…

Part II Outline

- Quant Strategies: Types, Parameters and Development
- Online Algorithms for Portfolio Optimization: Theory & Practice

Types of Quant Strategies

- Technical trading
  - signals for individual stocks based on price and volume history
  - examples: breakouts, moving average crossovers
  - also used as aids to understanding for human traders (the “chartists”)
- Pairs trading
  - bet on convergence of “related” stocks (e.g. Coke vs. Pepsi)
  - market-neutral
- Statistical modeling
  - regress stock on overall market returns, sector returns, other factors
  - wait for large deviations between model and empirical returns
  - PCA generalizations of pairs trading
- Event-driven
  - e.g. buy or sell a stock when analysts upgrade/downgrade
  - may be self-sizing
- Many signals have both a momentum and mean-reversion interpretation

Parameters for Optimization

- Universe of stocks
  - e.g. SP500, R2000, mid-caps, specific sectors, other criteria…
  - need to be very careful here…
- Timescales: trading frequency and holding period
  - constrains execution parameters
- Hedging method
  - reducing exposure: which indices to “subtract off”? (and there are many)
  - futures vs. options
- Strategy-specific parameters
  - thresholds for trading or length of position list
  - any parameters of the stat model
- Risk-return tradeoff
  - larger PNL vs. lower variance
- Trade execution
  - T and V from Part I (stock-specific)
  - method: VWAP, implementation shortfall, market on close,…
- Even “simple” ideas require a great deal of engineering

Strategy Development Process

- 1. A plausible high-level idea
  - “let’s buy/sell when analysts upgrade/downgrade”
  - “let’s apply Exponential Gradient to the long & short SP500”
- 2. Quick-and dirty backtesting
  - usually make crude/optimistic assumptions about execution costs, market impact, etc.
  - e.g. assume we can get market on close +/- round-trip bid/ask estimate
  - optimize strategy parameters
  - may not be possible for high-frequency intraday strategies
- 3. Evaluate performance
  - profitability and risk
  - scalability!
- 4. Get serious
  - improved realism in backtests
  - optimize execution parameters
  - explore various hedging methods
  - analyze exposures
  - make sure you understand why it works (and is “different”)
- 5. Cross fingers and begin live trading
  - usually at reduced volume initially
- 6. Monitor performance continually; adjust and resize

Online Algorithms for Portfolio Optimization: Theory and Practice

[Thanks to E. Even-Dar, C. Ural, J. Wortman]

Basic Framework

- An underlying universe of K assets U = {S_1,…,S_K}
- Goal: manage a "profitable" portfolio over U
  - each trading period S_i grows/shrinks q_i = (1+r_i), r_i in [-1,infinity]
  - we maintain a distribution w of wealth, fraction w_i in S_i
  - all quantities indexed by time t
- Traditionally: K assets are long positions in common stocks
- Generalized: K assets are any collection of investment instruments:
  - long and short positions in common stocks, cash, futures, derivatives
  - technical trading strategies, pairs strategies, etc. (search keywords?)
  - generally need instruments to be "stateless": can be entered at any time
- How do we measure performance relative to U?
  - average return (~"the market"): place 1/K of initial wealth in each S_i and leave it there
  - Uniform Constant Rebalanced Portfolio (UCRP): set w_i = 1/K and rebalance every period
  - Best Single Stock (BSS) in hindsight
  - Best Constant Rebalanced Portfolio (BCRP) in hindsight
  - Note: must place some restrictions on comparison class
- What about risk?
  - Sharpe Ratio = (mean of returns)/(standard deviation of returns)
  - Mean-Variance (MV) criterion = mean – variance
  - Maximum Drawdown
  - Value at Risk (VaR)
  - more refined: distinguishing "good" vs. "bad" variance

Online Algorithms: Theory

- Assume nothing about sequence of returns r_i (except maybe max loss)
- On arbitrary sequence r^1,…r^T, algorithm A dynamically adjusts portfolio w^1,…,w^t
- Compare cumulative return of BSS in hindsight to return of A
- Powerful family of no-regret algorithms: for all sequences,
  - Return(A(r^1,…,r^T))/T >= Return(BSS(r^1,…,r^T))/T – O(sqrt(log(K)/T))
  - per-step regret is vanishing with T
- How is this possible?
  - note: for this to be interesting, need BSS to strongly outperform the average
- Turns out to be crucial to update weights multiplicatively, not additively
- Flavor of a typical algorithm:
  - w_i $\leftarrow$ exp($\eta * $r_i)w_i, renormalize
- One (crucial) parameter: learning rate η
  - for the theory, need to optimize η ~ 1/sqrt(T)
  - generally are assuming momentum rather than mean reversion
  - note: η = 0 (no learning) is UCRP; a form of mean reversion
  - value of η also strongly influences portfolio concentration $\rightarrow$ variance/risk
- Let’s look at some empirical performance

Data Period: 1/4/1999 – 8/2/2005
Underlying Instruments: 466 stocks in S&P 500
Daily (closing) returns and trading
Mark-to-market

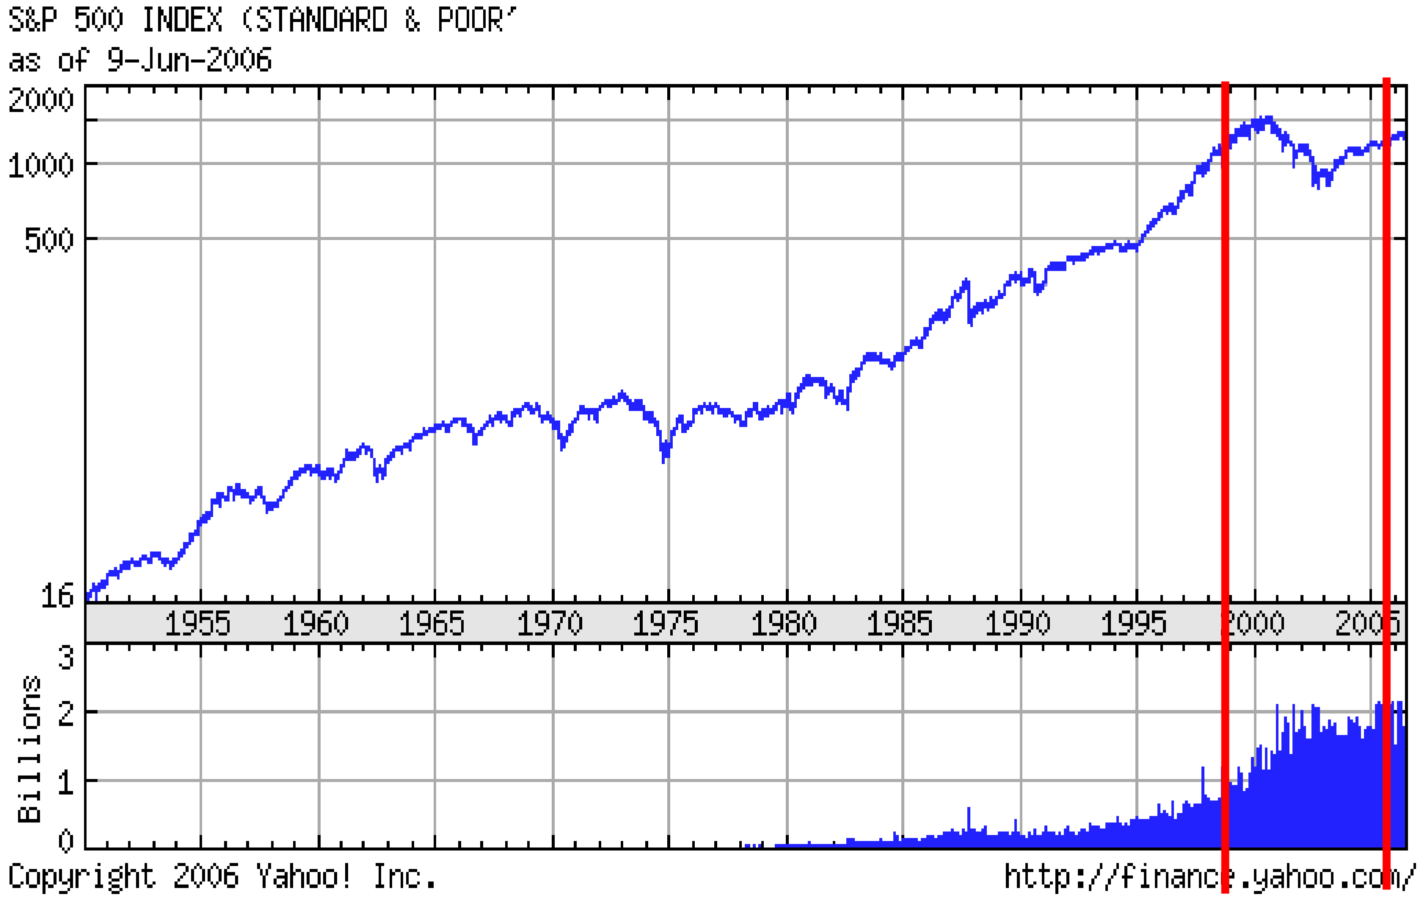

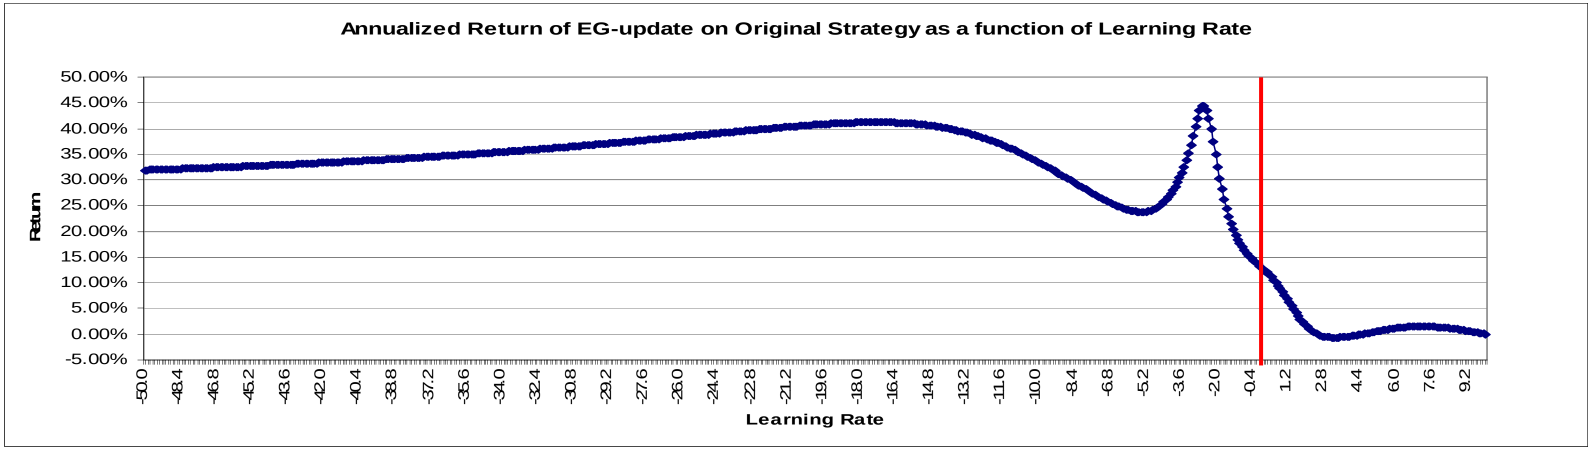

Annualized Return of EG-update on Original Strategy as a function of Learning Rate
Return
Learning Rate
50.00%
45.00%
40.00%
35.00%
30.00%
25.00%
20.00%
15.00%
10.00%
5.00%
0.00%
-5.00%
-50.0
-48.4
-46.8
-45.2
-43.6
-42.0
-40.4
-38.8
-37.2
-35.6
-34.0
-32.4
-30.8
-29.2
-27.6
-26.0
-24.4
-22.8
-21.2
-19.6
-18.0
-16.4
-14.8
-13.2
-11.6
-10.0
-8.4
-6.8
-5.2
-3.6
-2.0
-0.4
1.2
2.8
4.4
6.0
7.6
9.2

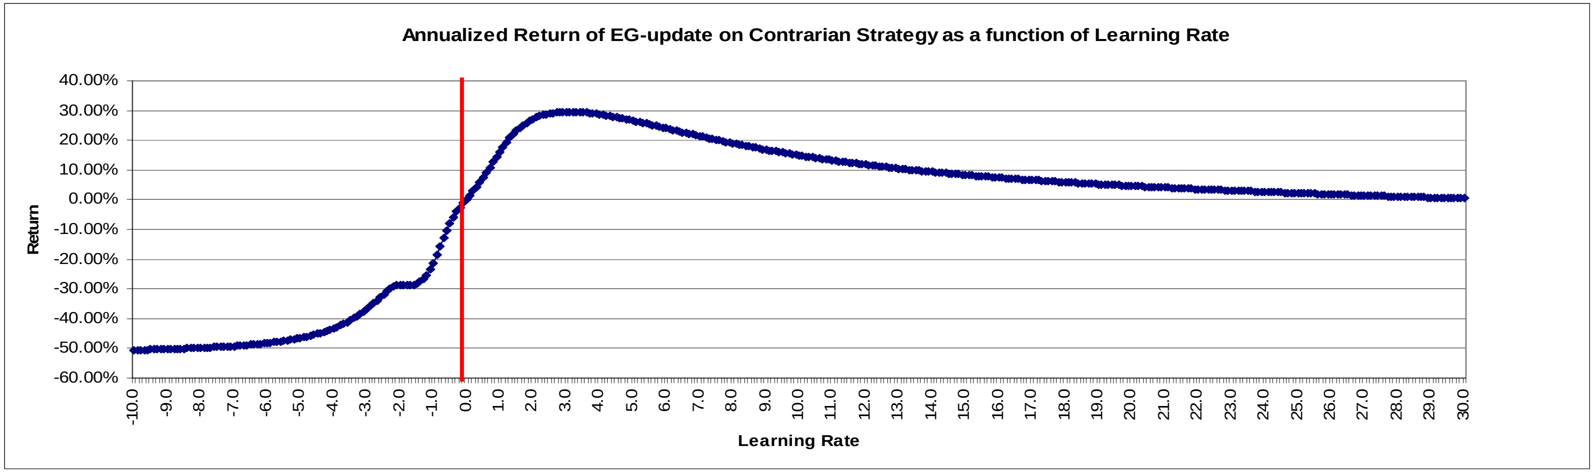

Annualized Return of EG-update on Contrarian Strategy as a function of Learning Rate
Return
Learning Rate
40.00%
30.00%
20.00%
10.00%
0.00%
-10.00%
-20.00%
-30.00%
-40.00%
-50.00%
-60.00%
-10.0
-9.0
-8.0
-7.0
-6.0
-5.0
-4.0
-3.0
-2.0
-1.0
0.0
1.0
2.0
3.0
4.0
5.0
6.0
7.0
8.0
9.0
10.0
11.0
12.0
13.0
14.0
15.0
16.0
17.0
18.0
19.0
20.0
21.0
22.0
23.0
24.0
25.0
26.0
27.0
28.0
29.0
30.0

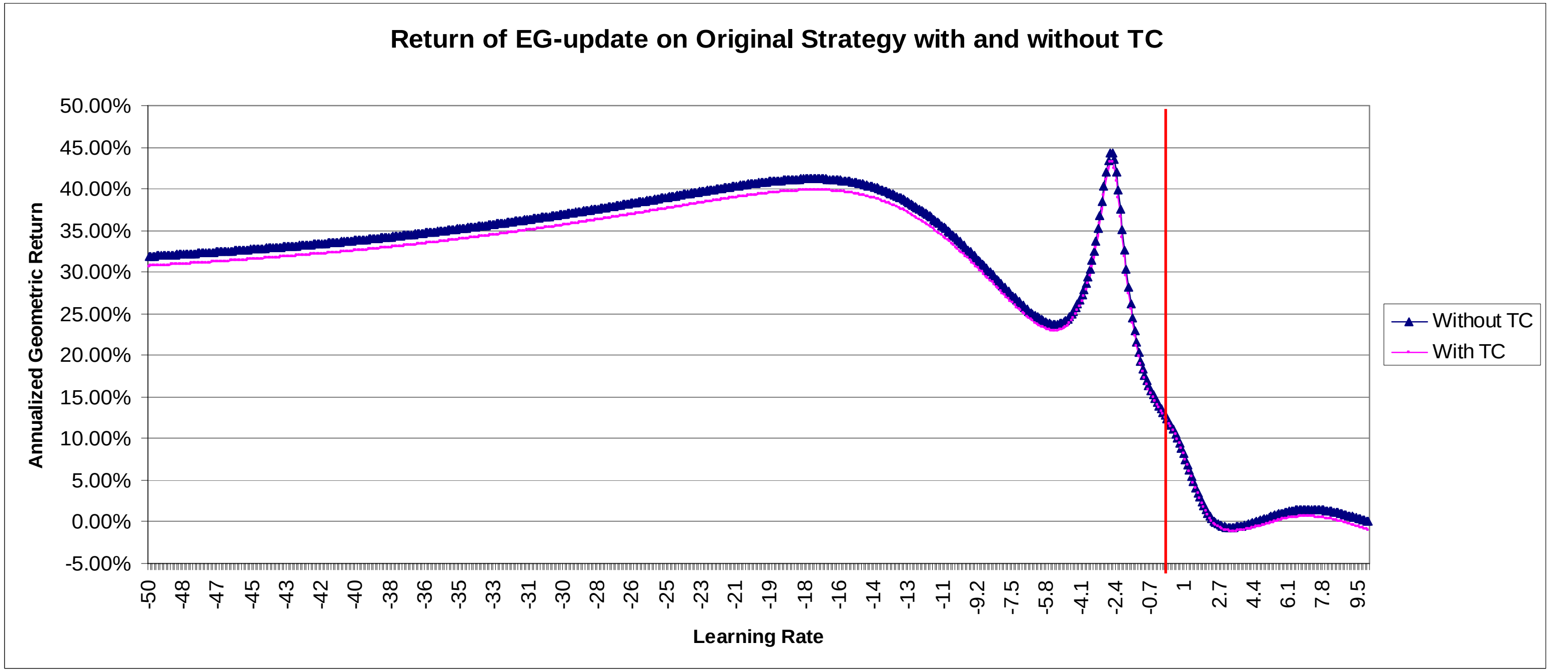

- Current Assumption: 4 bp.
- Not a significant effect to the Original Strategies (long only)

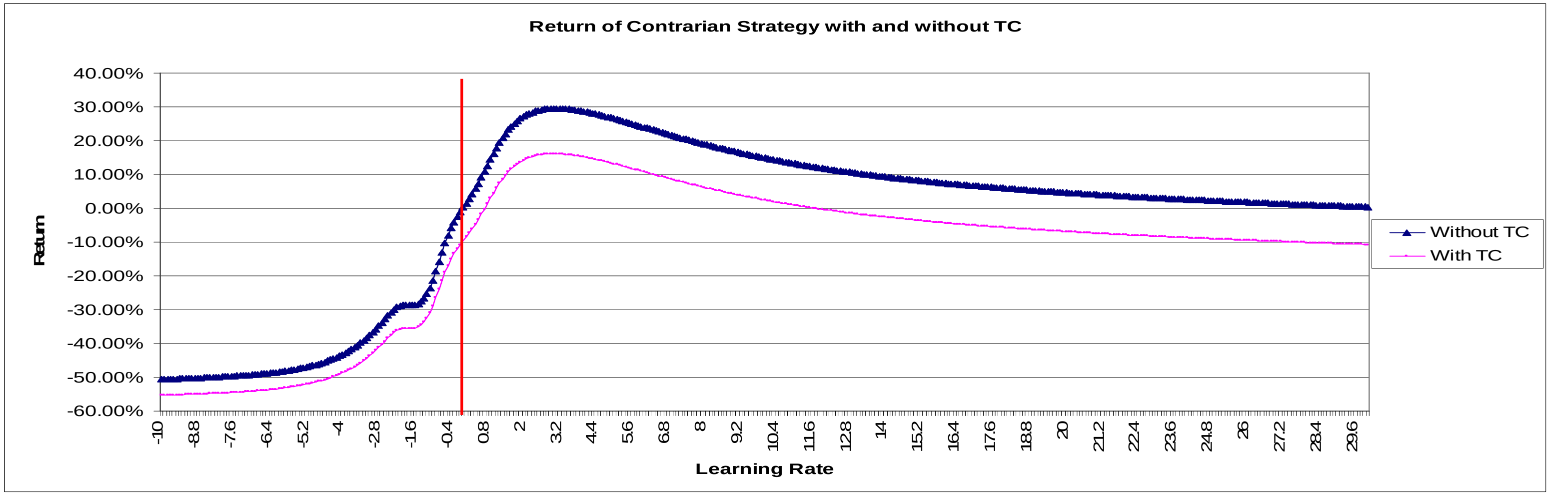
Return of Contrarian Strategy with and without TC
Return
Learning Rate
Without TC
With TC

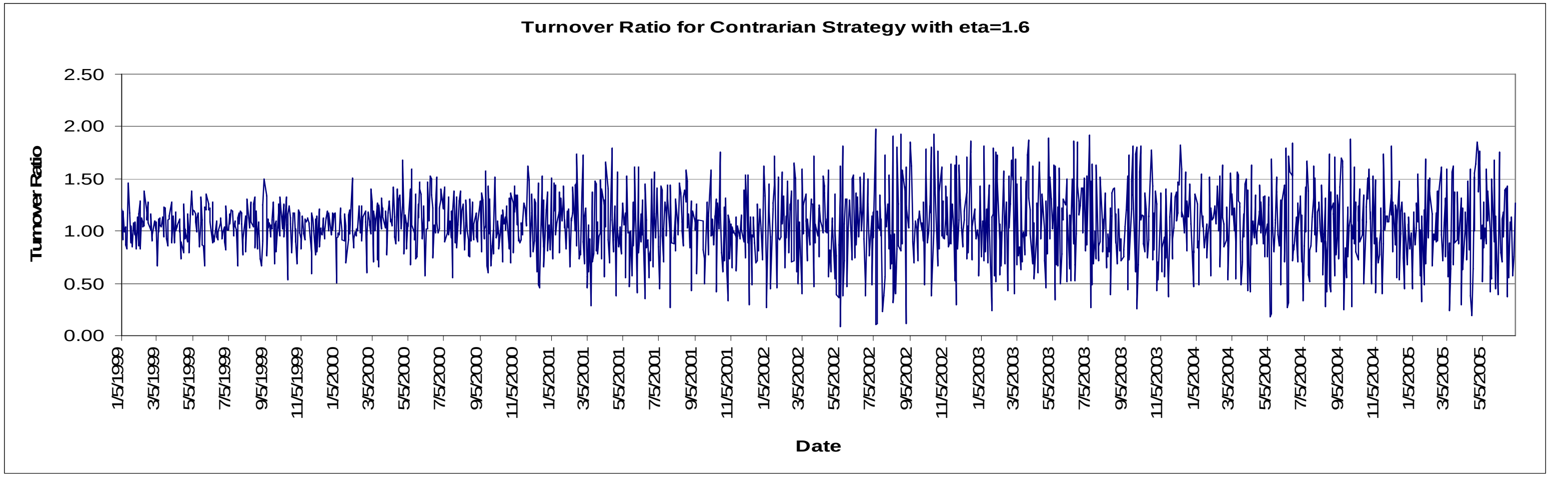
Turnover Ratio for Contrarian Strategy with eta=1.6
Turnover Ratio
Date

Invest $1 million (GMV) in the algorithm every day

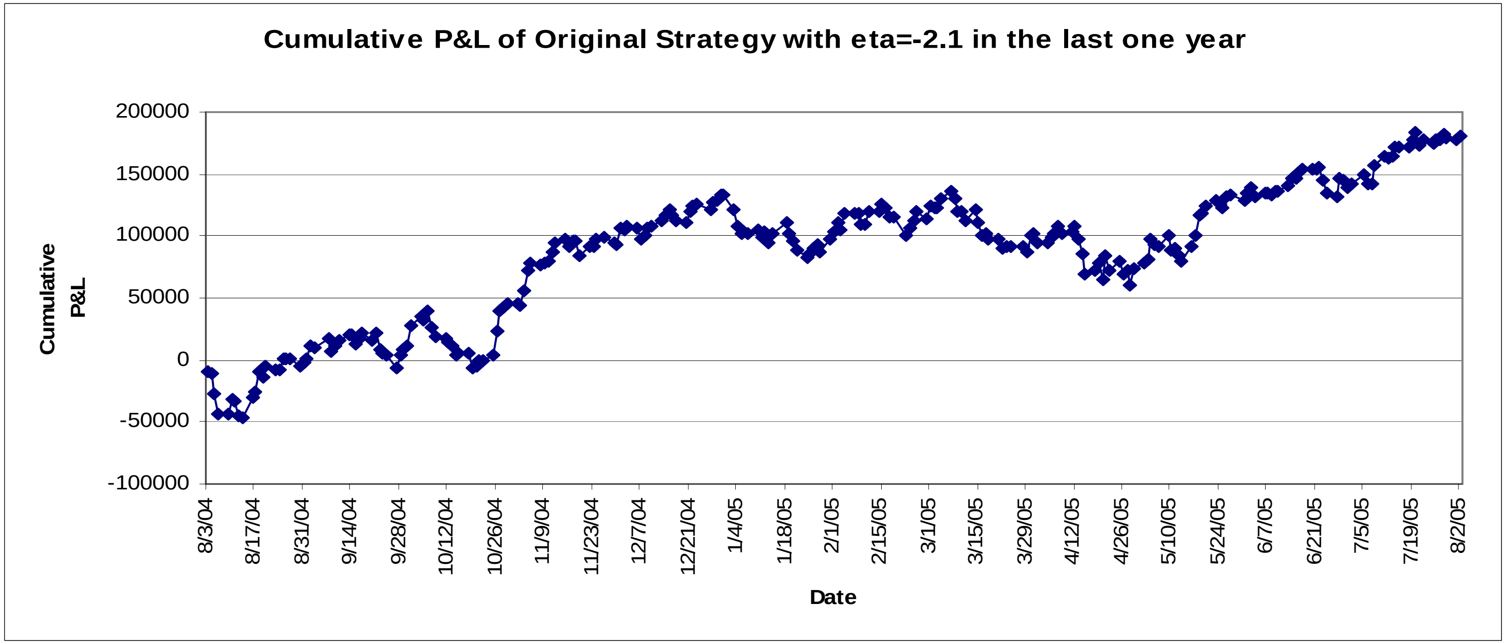

	EG (withTC)	S&P 500 Index
Annualized Arithmetic Return	17.94%	12.18%
Annualized Geometric Return	18.88%	12.32%
Annualized Stdev	11.31%	10.62%
Sharpe Ratio (Arithmetic)	1.59	1.15
Sharpe Ratio (Geometric)	1.67	1.16

What About Risk?

- Sharpe Ratio = (mean of returns)/(standard deviation of returns)
- Mean-Variance (MV) criterion = mean – variance
- Maximum Drawdown
- Value at Risk (VaR)
- More refined: distinguishing “good” vs. “bad” variance
- One (theoretical) ideal: no regret compared to BSS in hindsight w.r.t. risk-return
  - e.g. BSS Sharpe, BSS MV,…
  - can prove any online algorithm must have constant regret!
  - nevertheless…

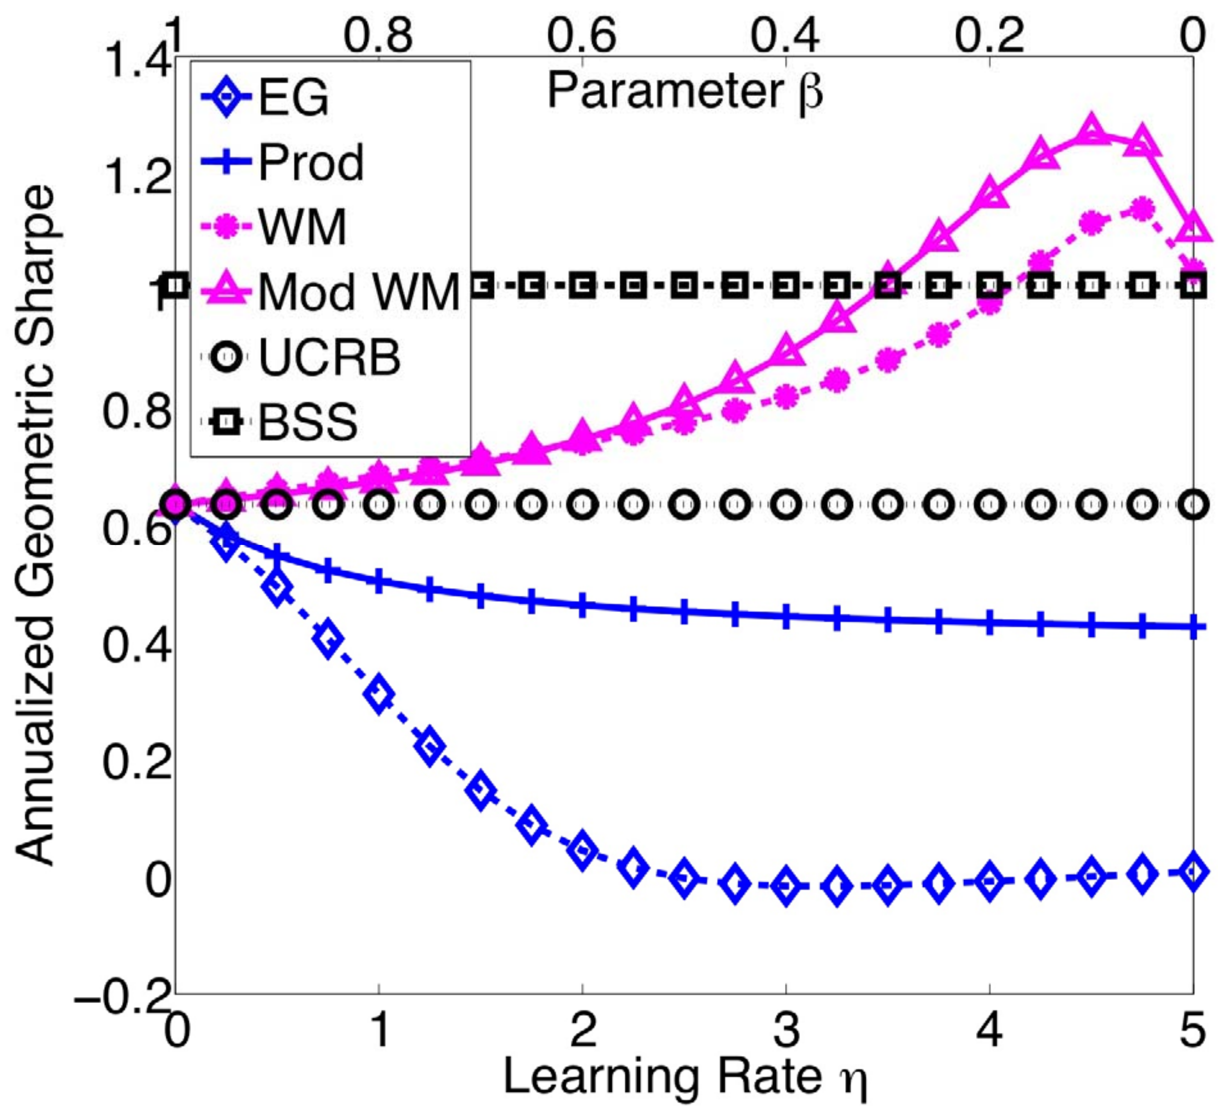

Parameter β
1
0.8
0.6
0.4
0.2
0
EG
Prod
WM
Mod WM
UCRB
BSS
Annualized Geometric Sharpe
1.4
1.2
1
0.8
0.6
0.4
0.2
0
−0.2
0
1
2
3
4
5
Learning Rate η

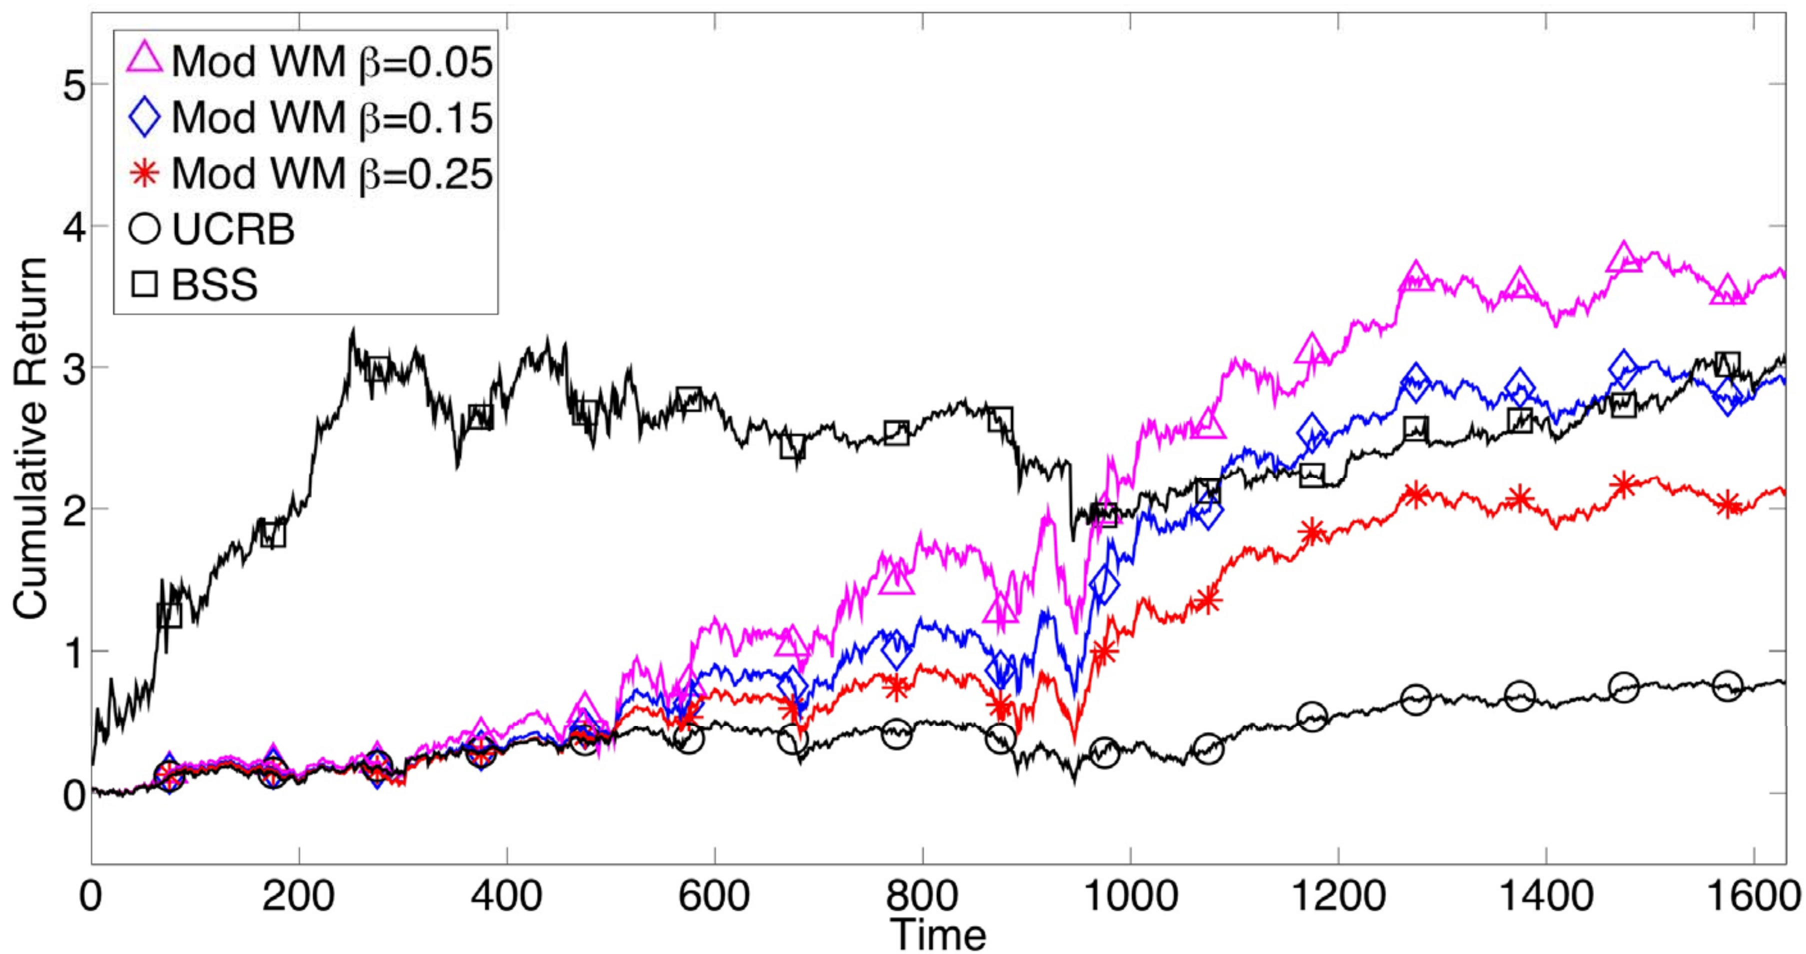

Mod WM β=0.05
Mod WM β=0.15
Mod WM β=0.25
UCRB
BSS
Cumulative Return
Time
0
1
2
3
4
5
0
200
400
600
800
1000
1200
1400
1600

Conclusions

- Many algorithmic challenges in modern finance
- Low level: optimized execution & microstructure
- High level: quant strategy design and development
- Space in between filling rapidly
- More speculative comments:
  - importation of finance methodology into emerging markets (search keywords)
  - the Optimark story