Algorithmic Trading and Machine Learning

Michael Kearns
University of Pennsylvania

QuantCon 2015, NYC

Special thanks: Yuriy Nevmyvaka (Lehman, BofA, SAC, Engineers Gate)

Long Version: http://techtalks.tv/talks/algorithmic-trading-and-machine-learning/60847/
or Google “techtalks ICML 2014 kearns”

ML for Trading: Challenges

- Learning to Act (vs. Predict): Optimized Execution
- Dealing with Censored Data: Order Routing in Dark Pools
- Incorporating Risk: Trading Under Inventory Constraints

Learning to Act: ML for Optimized Execution

[Y. Nevmyvaka. Y. Feng, MK; ICML 2006]
[MK, Y. Nevmyvaka; In “High Frequency Trading”, O’Hara et al. eds, Risk Books 2013]

A Canonical Trading Problem

- Goal: Sell V shares in T time steps; maximize revenue
- Benchmarks:
  - Volume Weighted Average Price (VWAP)
  - Time Weighted Average Price (TWAP)
  - Implementation Shortfall (midpoint of bid-ask spread at beginning)
- View as a problem of *state-based control (Reinforcement Learning)*
  - Action space: limit orders
  - State variables: inventory and time remaining
  - Additional features capturing order book activity
- Experimental framework
  - Full historical order book reconstruction and simulation
  - Learn optimal policy on 1 year training; test on following 6 months
  - Pitfalls: directional drift, “counterfactual” market impact

LAST MATCH		TODAY'S ACTIVITY	
Price	23.7790	Orders	1,630
Time	9:01:55.614	Volume	44,839

BUY ORDERS		SELL ORDERS	
SHARES	PRICE	SHARES	PRICE
1,000	23.7600	100	23.7800
3,087	23.7500	800	23.7990
200	23.7500	500	23.8000
100	23.7400	1,720	23.8070
1,720	23.7280	900	23.8190
2,000	23.7200	200	23.8500
1,000	23.7000	1,000	23.8500
100	23.7000	1,000	23.8500
100	23.7000	1,000	23.8600
800	23.6970	200	24.0000
500	23.6500	500	24.0000
3,000	23.6500	1,000	24.0300
4,300	23.6500	200	24.0300
2,000	23.6500	1,100	24.0400
200	23.6200	500	24.0500
(195 more)		(219 more)	

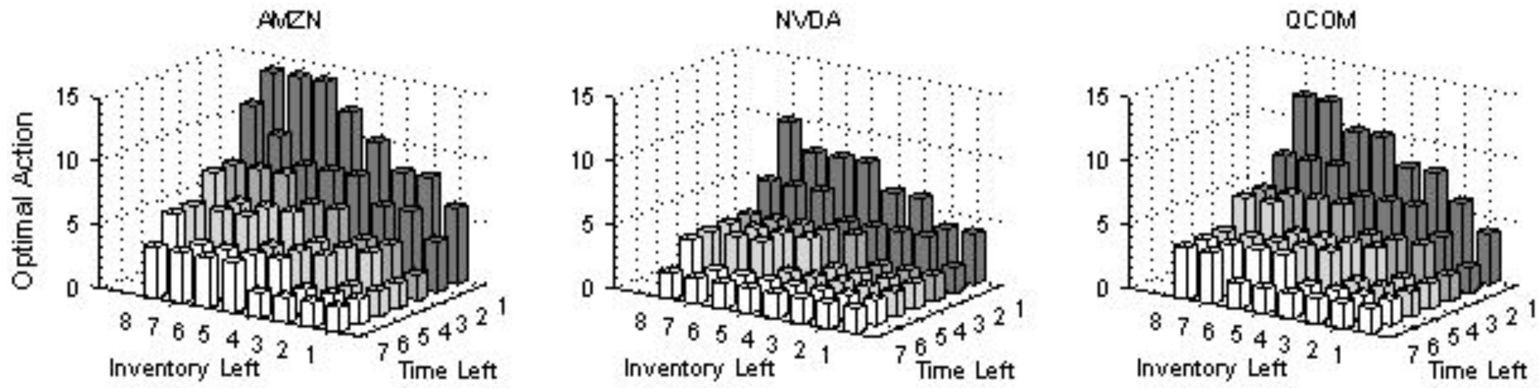

Improvement Over Optimized Submit-and-Leave

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%

Additional 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%**

Desperately Seeking Alpha

- A natural modification:
  - Change action space to buy or sell and hold for t seconds, then liquidate (+null action)
  - Add state features capturing directional movements
- Now trying to predict movement and *profit* (vs. fixed optimization problem)
- Definite (aggregate) predictability, but hard to overcome *trading costs*
- Still learn broadly consistent policies across stocks:
  - Null action vast majority of time; trade only in extremal states/opportunities
  - Short holding (milliseconds): Momentum
  - Longer holding (seconds): Reversion

Smart Order Routing in Dark Pools

[K. Ganchev, MK, Y. Nevmyvaka. J. Wortman Vaughan; UAI 2009, CACM 2010]
[K. Amin, MK, P. Key, A. Schwaighofer; UAI 2012]

Dark Pools

- Recently introduced trading mechanism
- Intended to allow large counterparties to trade with minimal market impact
- Only specify desired volume and direction (buy/sell); *no price specified*
- Buyers and sellers matched in order of arrival
- Prices will be midpoint of National Best Bid and Offer (NBBO) in *lit market*
- Now dozens of dark pool, competing for *liquidity* instead of price
- Break trade up over *exchanges* instead of over *time*

Smart Order Routing (SOR)

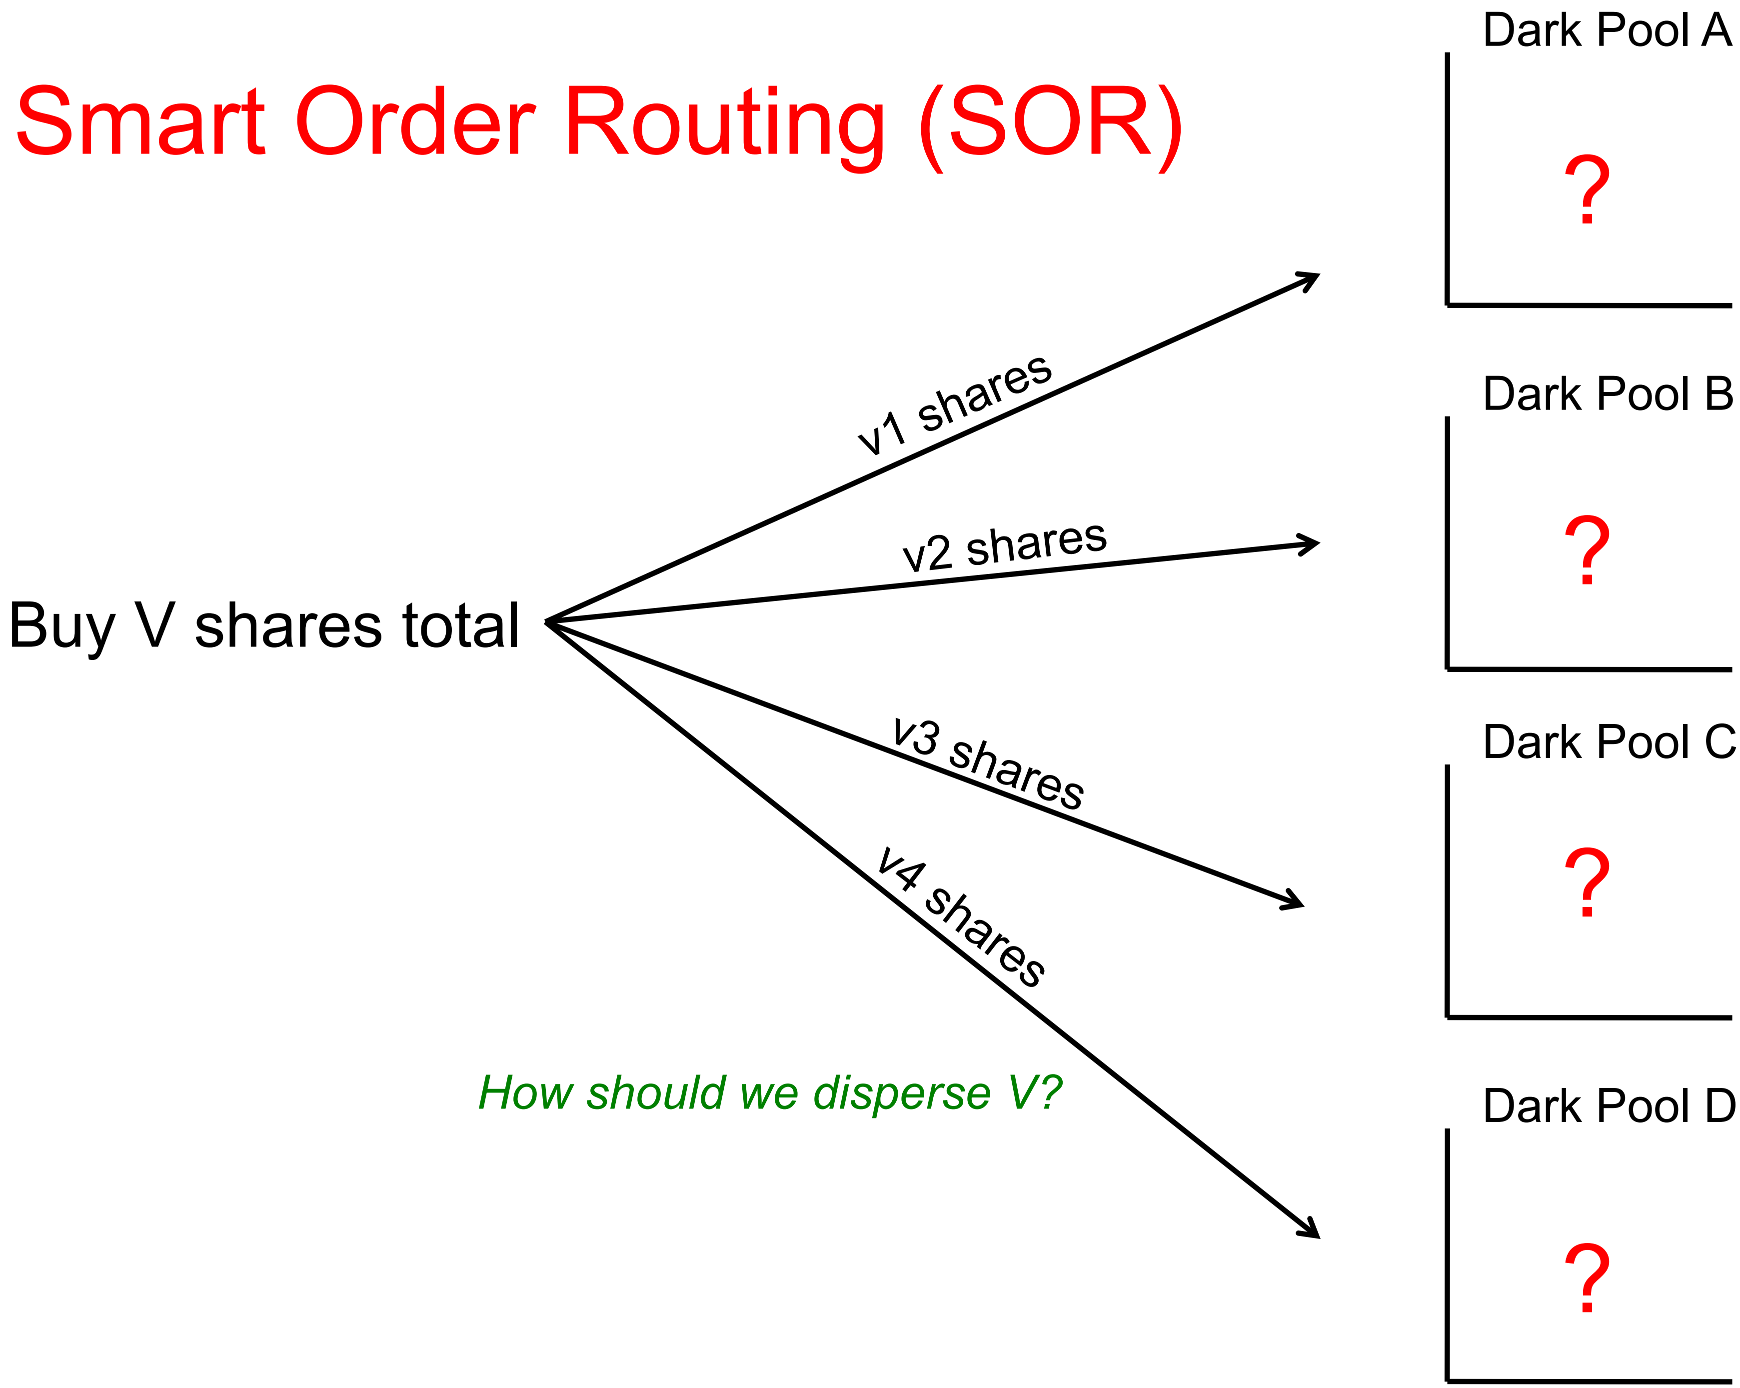

A Distributional Model of Liquidity

- Assume each dark pool has a stationary distribution P over available shares
- If we submit v shares, min(v,s) will be executed where s ~ P
- Our observations are *censored* by our own actions
- MLE for P is Kaplan-Meier --- but we must address *exploration* across pools
- Want to learn just enough about each pool to do optimal SOR

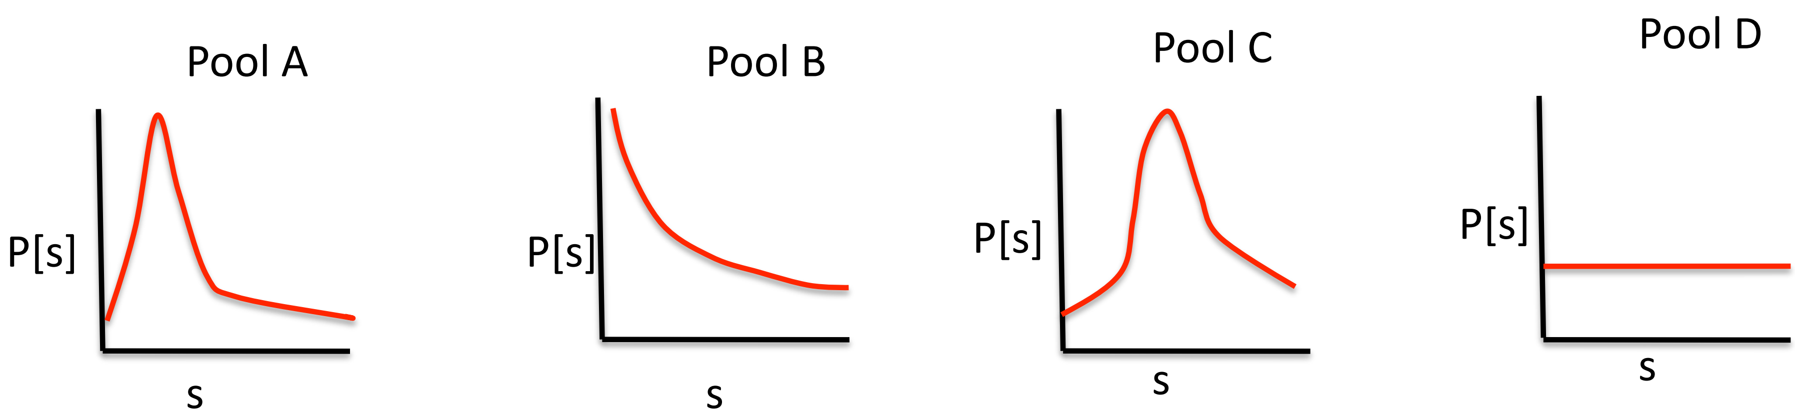

A Simple and Efficient Algorithm

Algorithm 2: Main algorithm.

Input: Volume sequence $V^1, V^2, V^3, \ldots$

Arbitrarily initialize $\hat{T}_i^1$ for each i;

for $t \leftarrow 1, 2, 3, \ldots$ **do**

    `% Allocation Step:`

    $\vec{v}^t \leftarrow \mathit{Greedy}(V^t, \hat{T}_1^t, \ldots, \hat{T}_K^t)$; greedy allocation under current distributional estimates

    for $i \in \{1, \ldots, K\}$ **do**

        Submit v_i^t units to venue i;

        Let r_i^t be the number of shares sold;

        `% Reestimation Step:`

        $\hat{T}_i^{t+1} \leftarrow \mathit{OptimisticKM}(\{(v_i^\tau, r_i^\tau)\}_{\tau=1}^t)$; re-estimate using censored observations

    end

end

- Provably converges quickly to optimal allocations under known distributions
- Involves optimistic modification to MLE, new convergence bound
- Analysis reminiscent of E^3/RMAX in RL

Empirical Evaluation

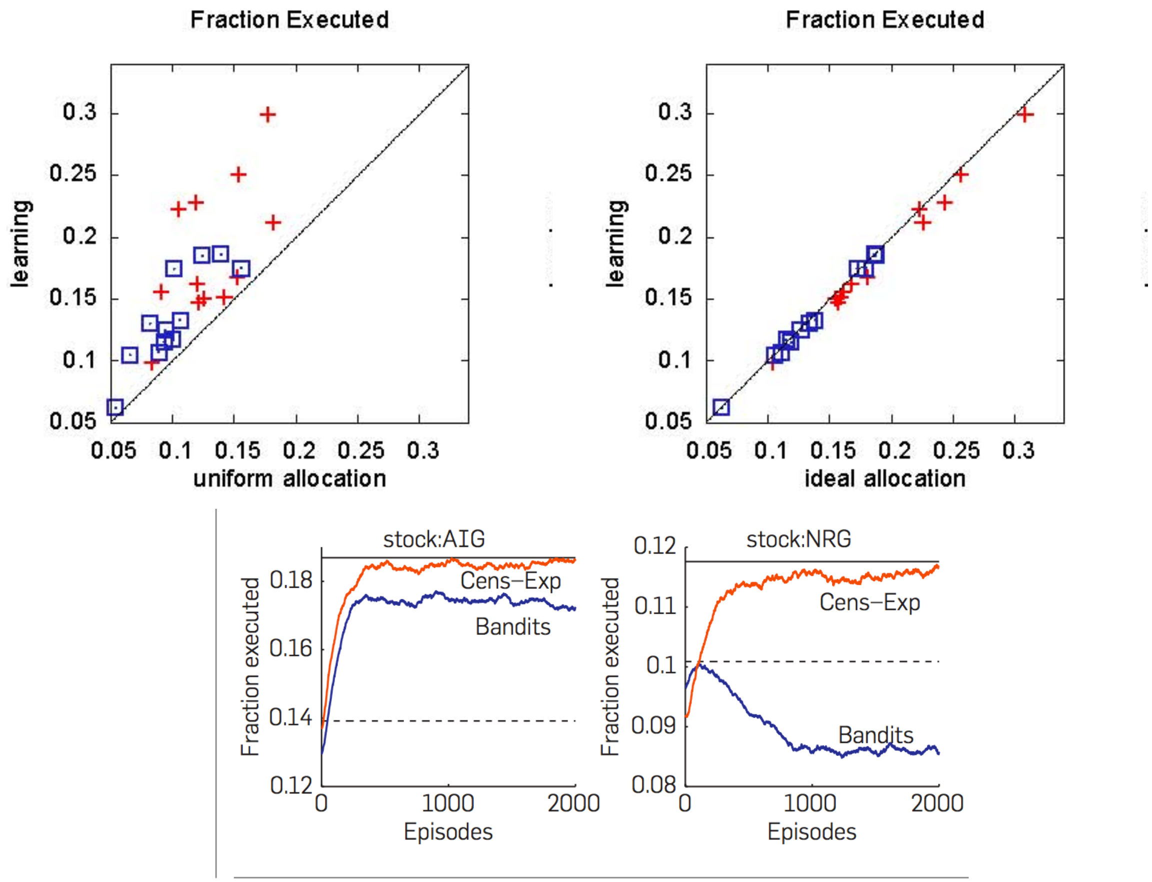

- Data: submission/execution data from multiple pools at large brokerage
- Used to build distribution models (heavy-tailed) and simulator
- Comparison to uniform allocation (strawman), bandit approach, optimal

Incorporating Risk: Algorithmic Trading with Inventory Constraints

[E. Even-Dar, MK, J. Wortman Vaughan; ALT 2006]
[L. Dworkin, MK, Y. Nevmyvaka; ICML 2014]

No-Regret Learning in Finance

- Originates with Cover’s Universal Portfolios; simple reweighting algorithm
- Strong theoretical guarantees *without stochastic assumptions*
  - Compete with best single stock in hindsight
- Can be applied directly to stocks or higher-level trading strategies
- Unfortunately methods work poorly in practice:

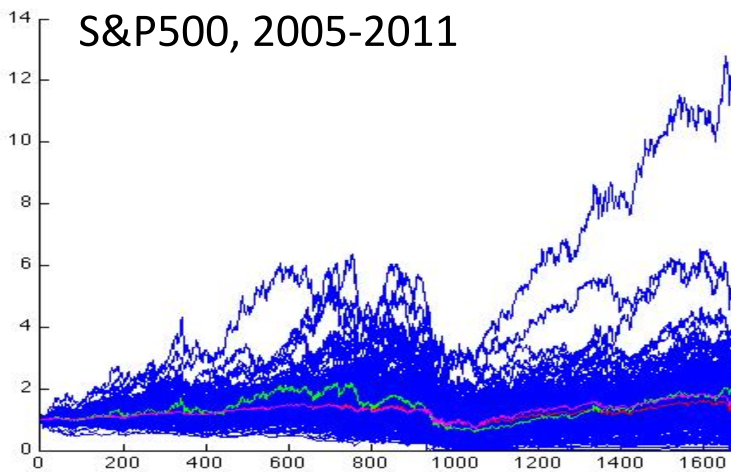

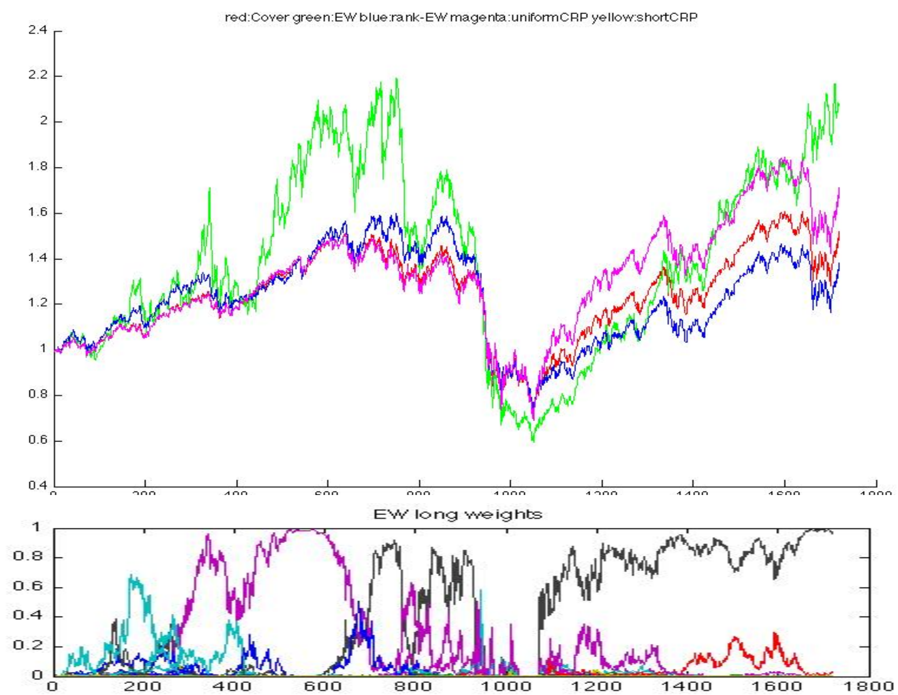

Trading with Inventory Constraints

- Can't manage to Sharpe Ratio, but can limit allowed positions/portfolios
- Restrict to portfolios with daily standard deviation PNL at most $X historically
- Leads to elliptical constraint in portfolio space depending on correlations
- Only compete with strategies:
  - Obeying inventory constraints
  - Making only local moves (limit market impact)
- Combine no-regret with pursuit-evasion to recover theoretical guarantees

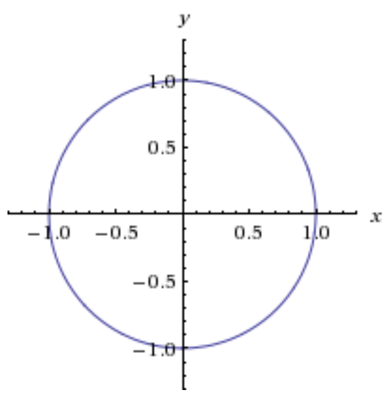

correlation = 0

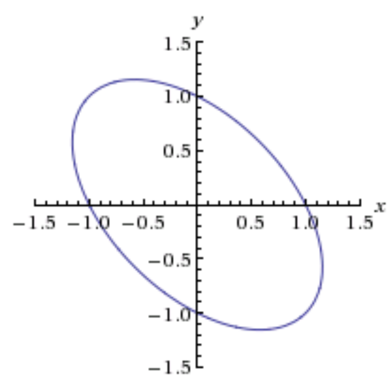

correlation = 0.5

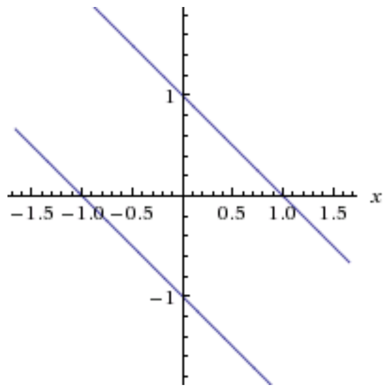

correlation = 1

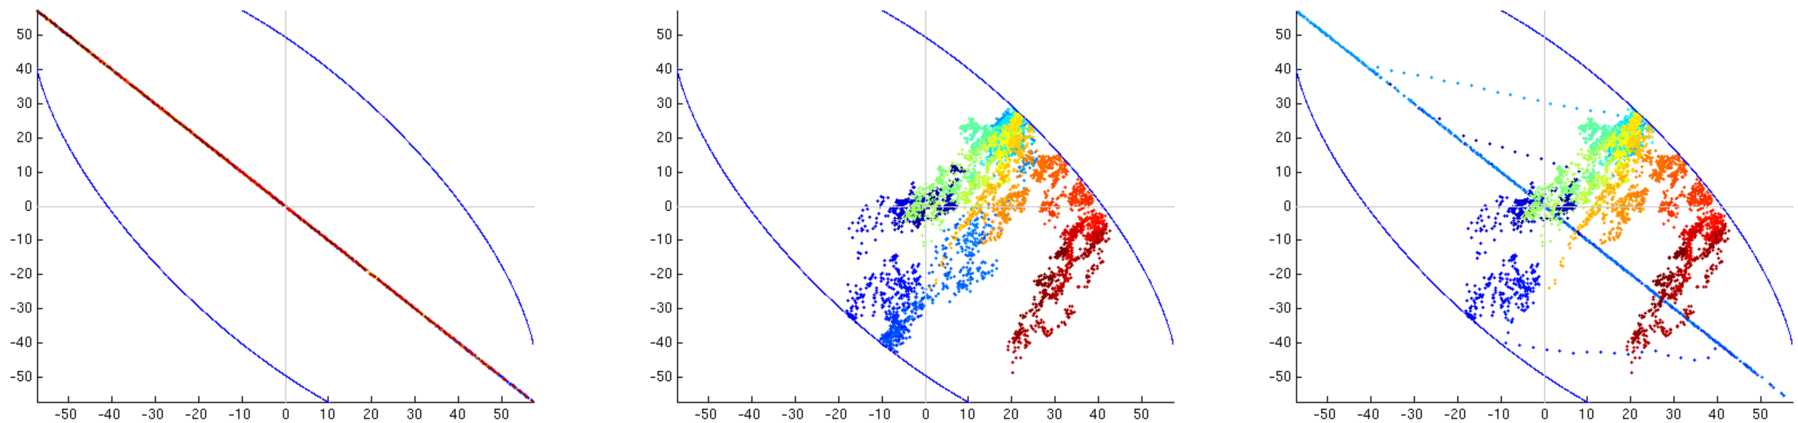

Hedged

Directional

Pursuit-Evasion

Conclusions

- In the middle (beginning?) of a period of rapid change in markets:
  - Automation of traditional processes and trading
  - Introduction of new market mechanisms (open order books, dark pools)
  - Development of new types of trading and strategies (HFT)
- Automation + Data → Machine Learning
- Challenges:
  - Feature design
  - Censored observations
  - Risk considerations
  - Strategic/adversarial behavior
- More, and different, to come…

Contact: mkearns@cis.upenn.edu

Long Version: http://techtalks.tv/talks/algorithmic-trading-and-machine-learning/60847/
or Google "techtalks ICML 2014 kearns"