Market Microstructure and Algorithmic Trading

PIMS Summer School 2016
University of Alberta, Edmonton
Lecture 4: July 6, 2016

Robert Almgren

quantitativebrokers

Trade trajectory determination

Overall schedule of trade

- Actual order placement by "micro" algo

Schedule to receive "best" price

- relative to benchmark
- minimize variance

Real implementation different than theory

- but theory is useful for ideas

Real examples

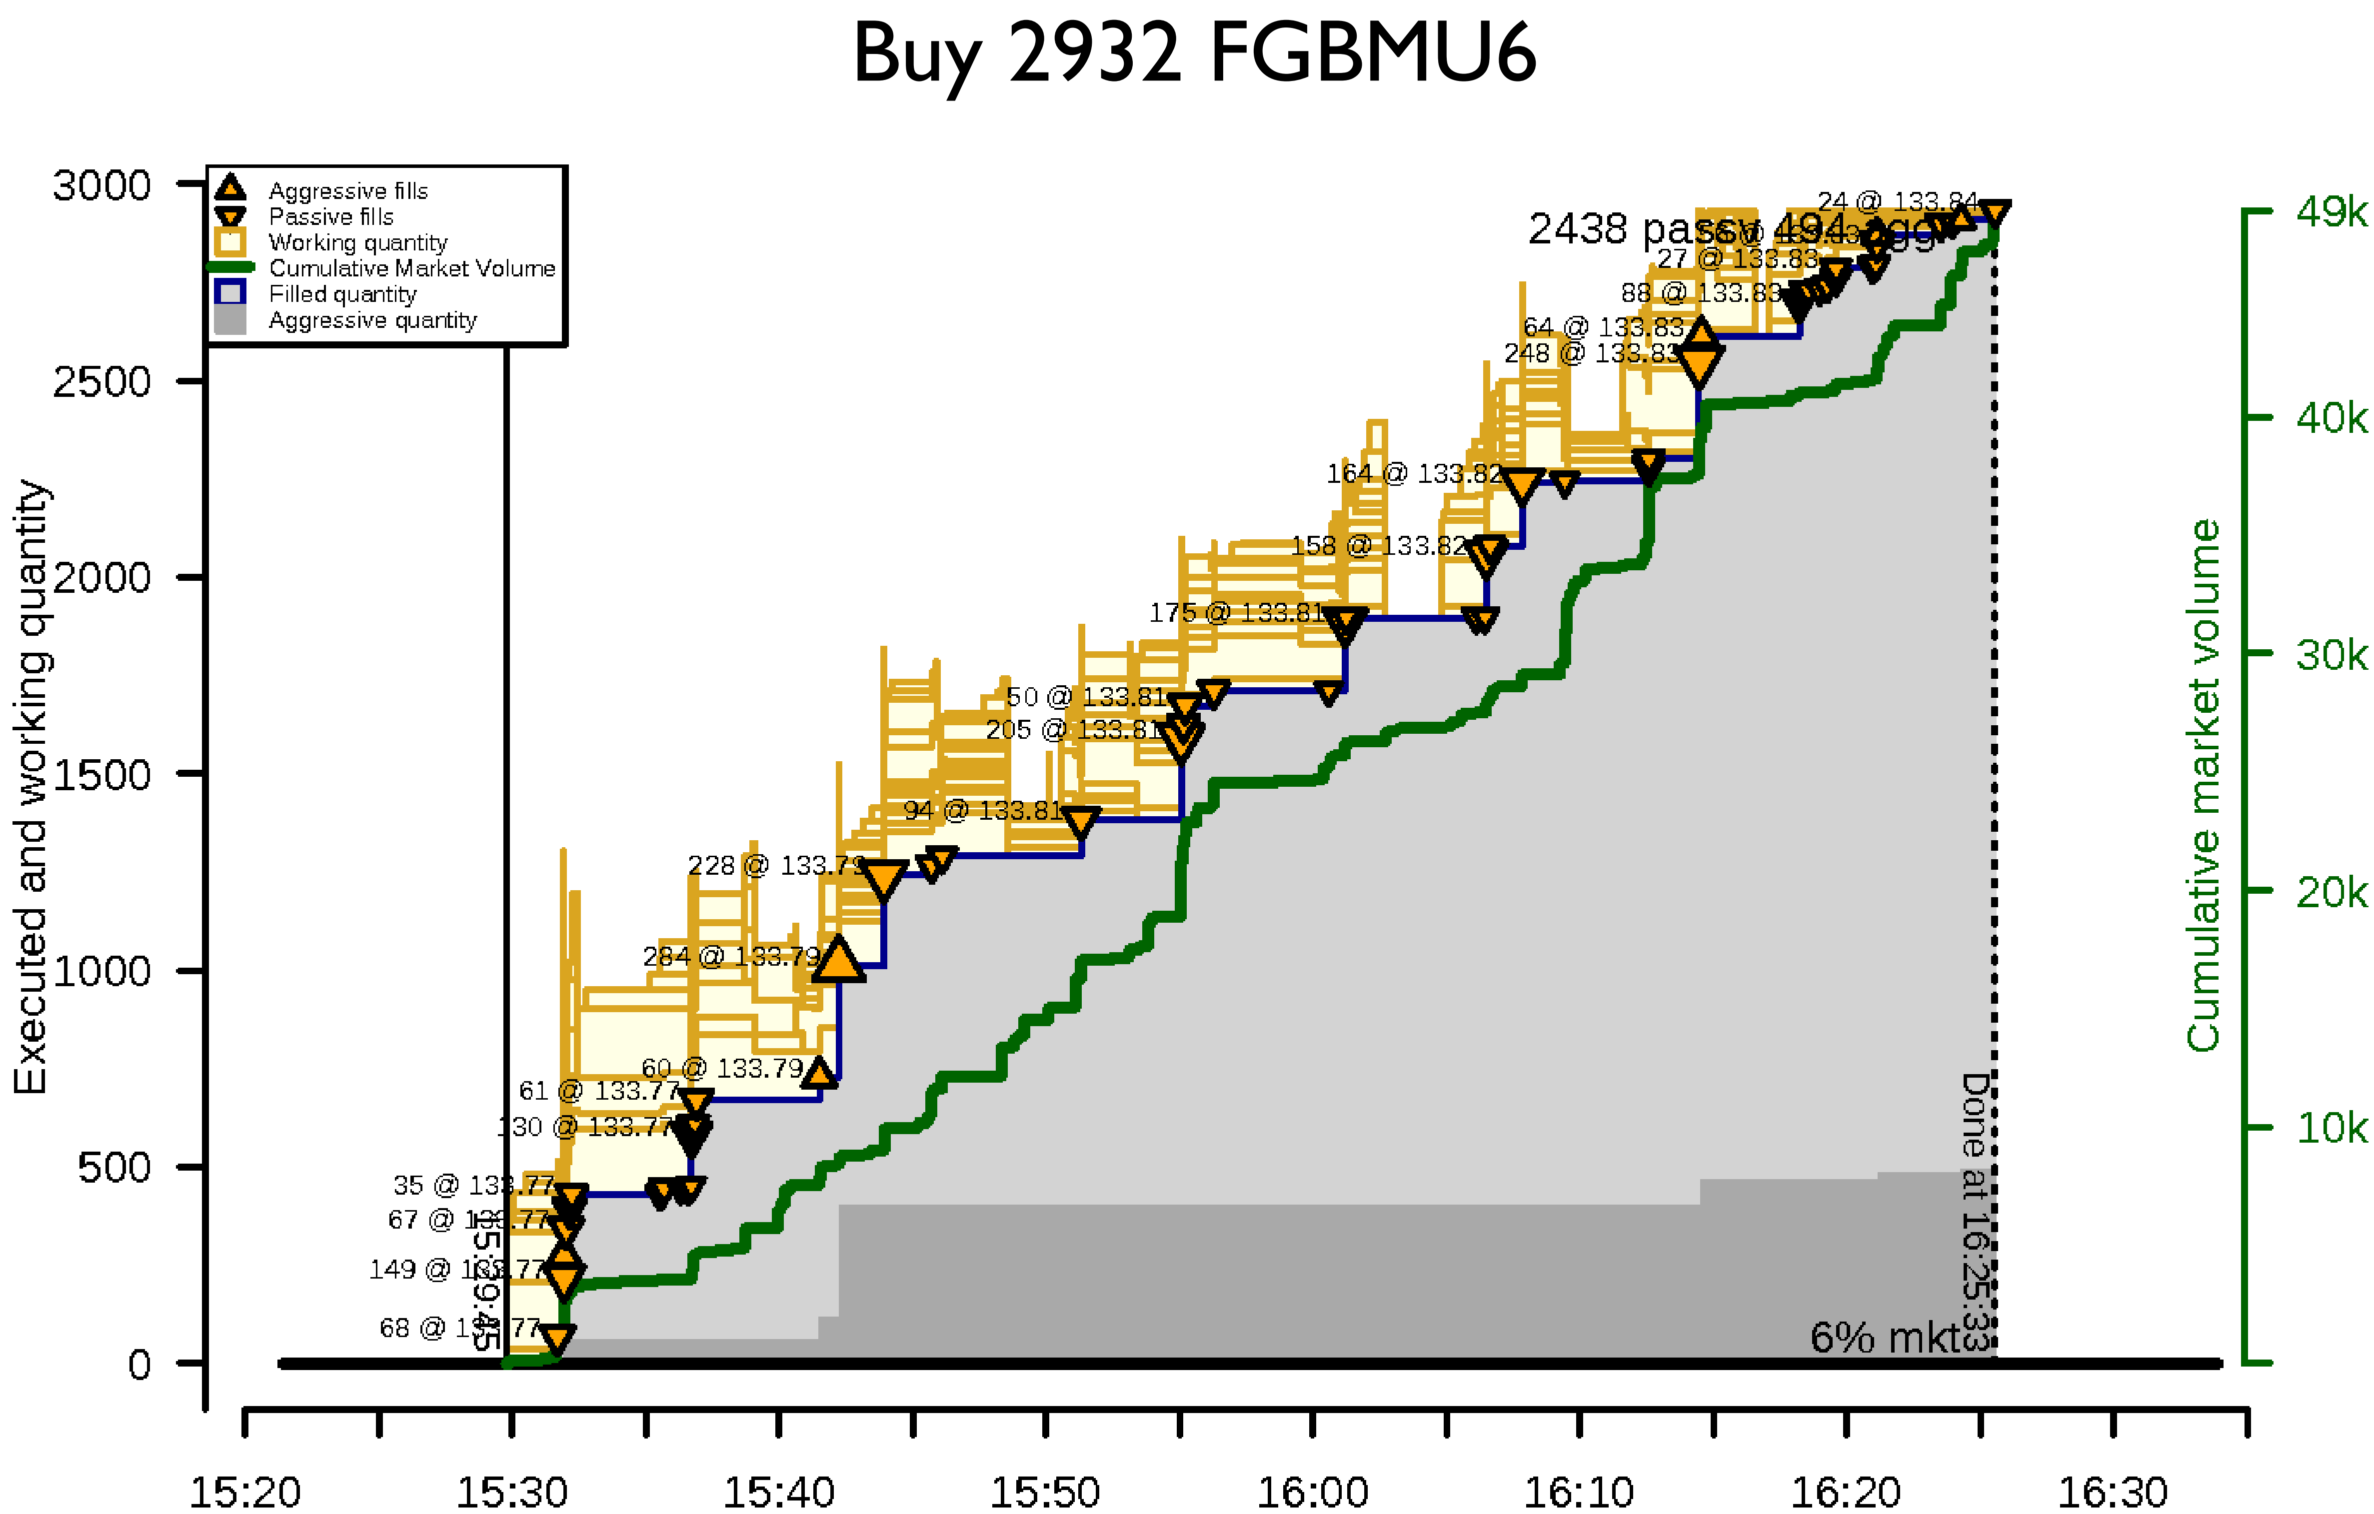

Produced by Quantitative Brokers from internal and Eurex Market Data

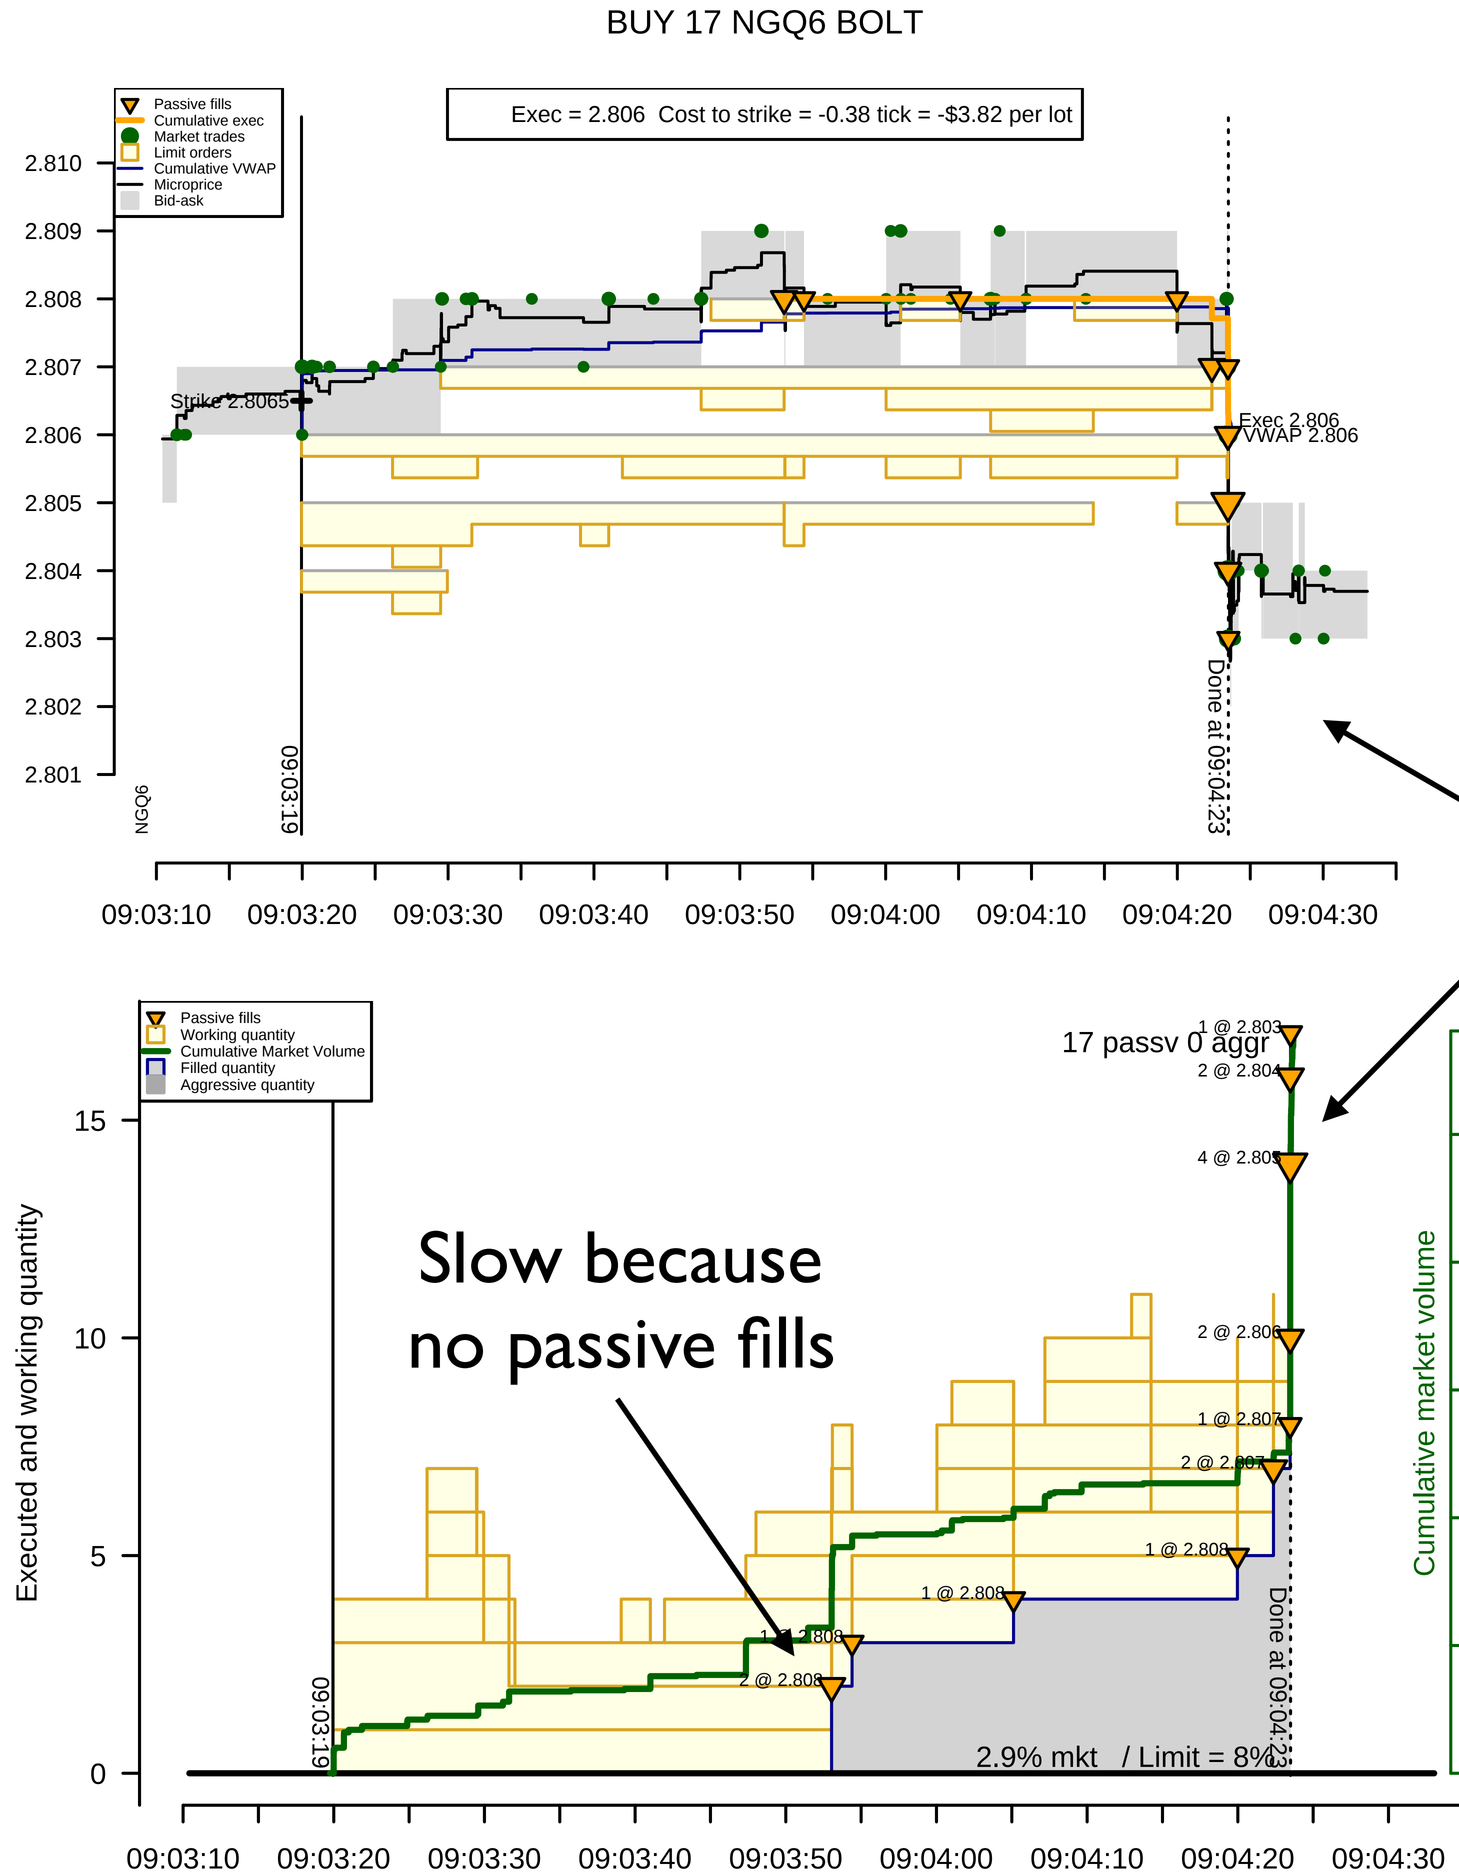

Sudden fills at end

Slow because no passive fills

Produced by Quantitative Brokers from internal and CME Market Data

Outline

1. Benchmarks
2. Market impact models
3. Mean-variance trajectory optimization
4. Option hedging

1. Benchmark

"Contract" with client

What would constitute an ideal trade

Choice of benchmark determines algo

Main benchmarks:

- Arrival Price
- VWAP
- TWAP

IS ≡ "Arrival Price"

The implementation shortfall: Paper versus reality

Reality involves the cost of trading and *the cost of not trading*.

André F. Perold

Journal of Portfolio Management; Spring 1988;
pg. 4

After selecting which stocks to buy and which to sell, "all" you have to do is implement your decisions. If you had the luxury of transacting on paper, your job would already be done. On paper, transactions occur by mere stroke of the pen. You can transact at all times in unlimited quantities with no price impact and free of all commissions. There are no doubts as to whether and at what price your order will be filled. If you could transact on paper, you would always be invested in your ideal portfolio.

To calculate the performance of the paper portfolio, you use the principle that on paper you transact instantly, costlessly, and in unlimited quantities. For example, if you would like to buy 50,000 shares at current prices, simply look at the current bid and ask, and consider the deal done at the average of the two. The same applies if you want to sell.

Using the average of the prevailing bid and ask means that you get the same price whether you are buying or selling. If you bought at the ask and sold at the bid, you would be incurring transaction costs. These occur only in real world implementations, not on paper.

If you trade quickly and aggressively, you will tend to pay a bigger price to transact. It is much harder to find the other side over the next hour than over the next week. When you are in a hurry, you also indicate your need to get in or out, which in turn may signal valuable information to others. Hence, the faster you trade, the larger your execution costs will be. On the other hand, you will have more of your ideal portfolio in place, and your opportunity costs consequently will be lower.

If you trade slowly and patiently, your execution costs will tend to be lower. For example, if you execute a large order in deliberate piecemeal fashion, you will not disturb the market very much. Alternatively, if you do not break up the order but bide your time until the other side shows up in size, then you may even reap a premium to market. Nevertheless, although your execution costs will be lower, your opportunity costs will be higher. For the more slowly you trade, the more you will be forgoing the fruits of your research, and the more you will become prone to adverse selection (which shows up mostly in opportunity cost). The longer you are out there, the more time others have to act strategically against you.

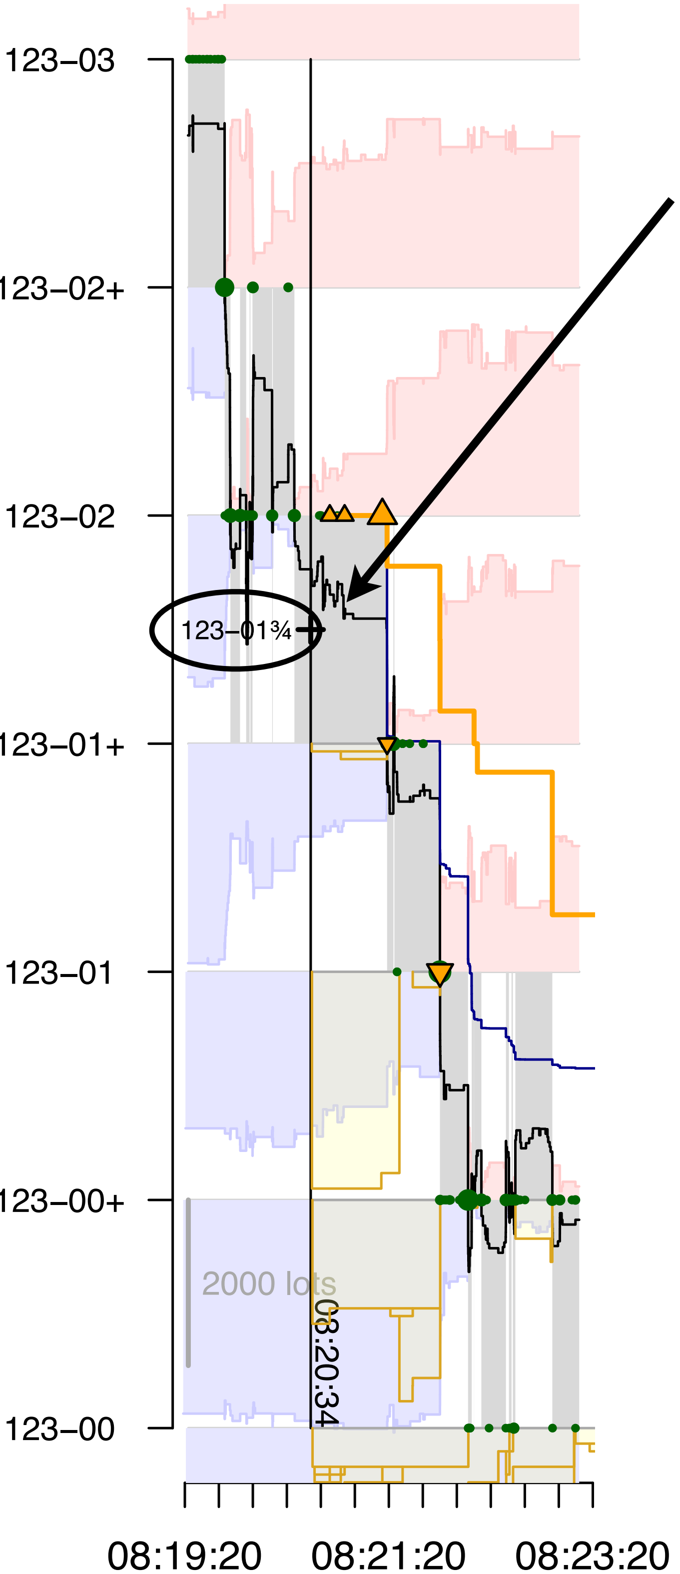

Arrival Price

Arrival price positives:

- corresponds to real trade decision
- cannot be gamed by executing broker

Arrival price negatives:

- subject to price motion
  - impossible to separate alpha from impact
- statistically very noisy

Minimize slippage to arrival price benchmark
by completing trade as rapidly as possible

Produced by Quantitative Brokers from internal and CME Market Data

Arrival Price Example

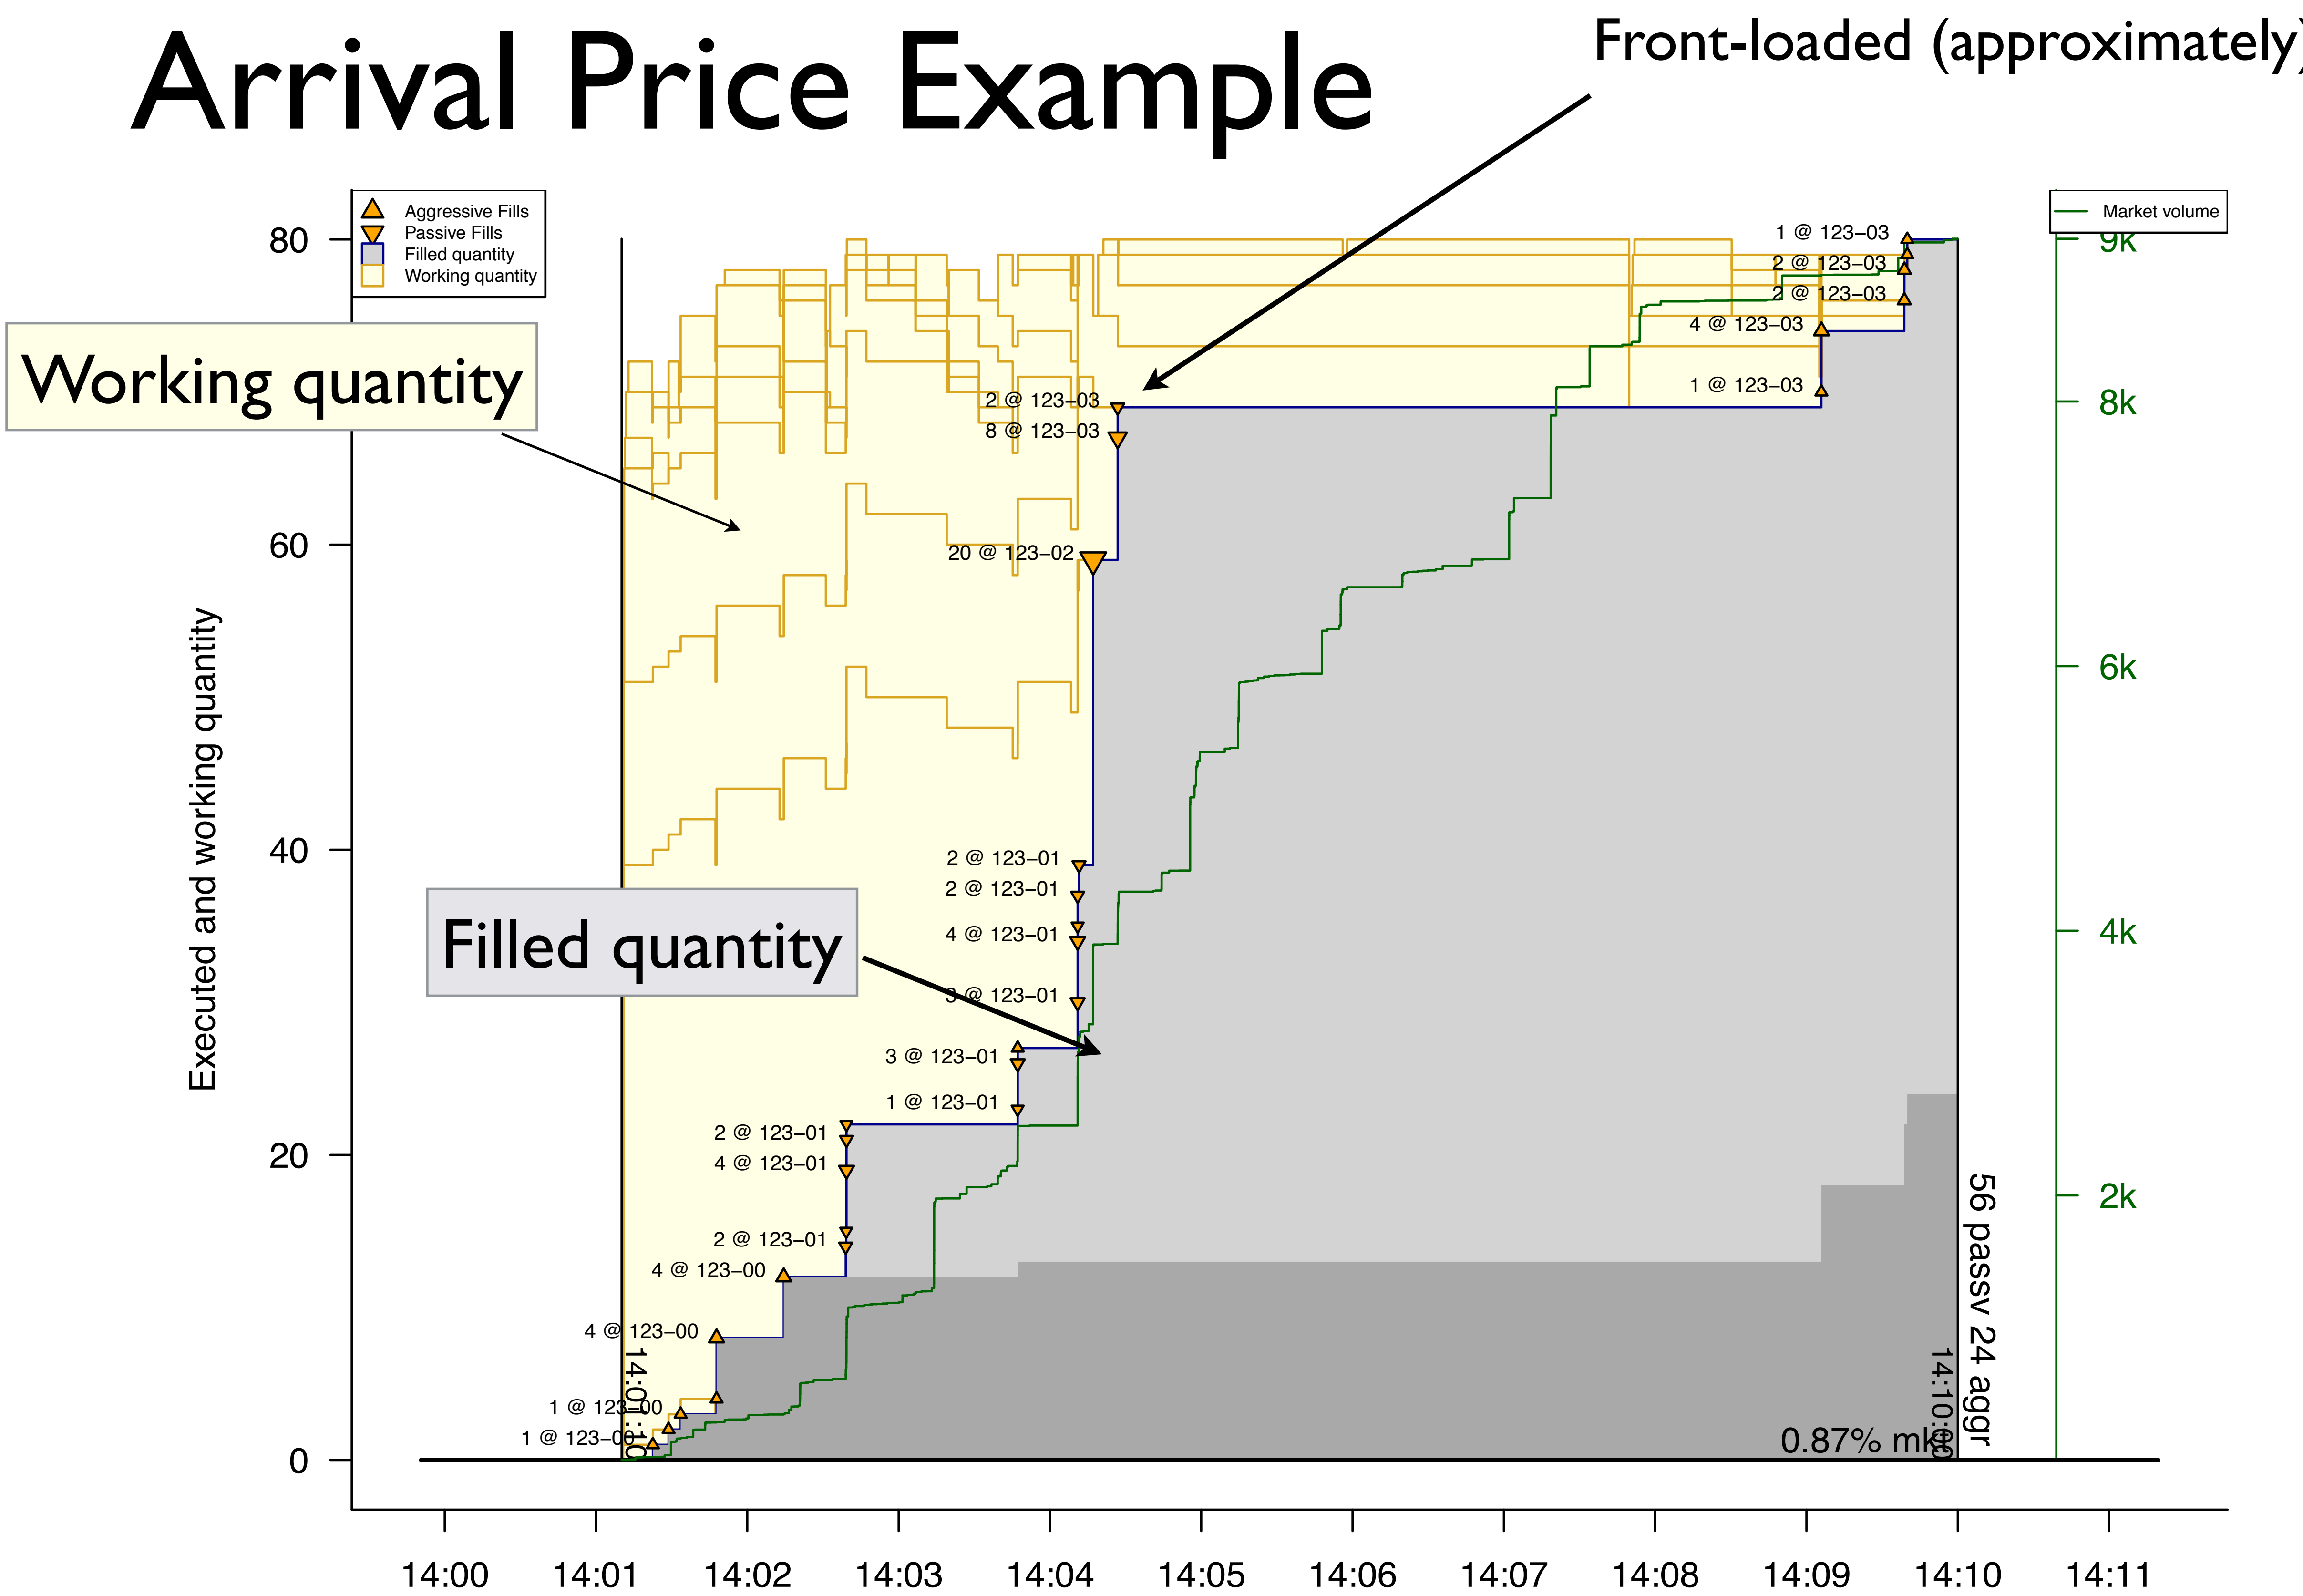

Produced by Quantitative Brokers from internal and CME Market Data

VWAP benchmark

Volume-Weighted Average Price

- between given start time and end time
- "what the market did"

Positives:

- easy to evaluate ex post: `select siz wavg prc`
- slippage is very consistent

Negatives:

- Does not correspond to any investment goal
- Can be gamed by broker
- Difficult to forecast volume profile during trading

Minimize slippage to VWAP benchmark
by trading exactly along market profile

VWAP Example (1)

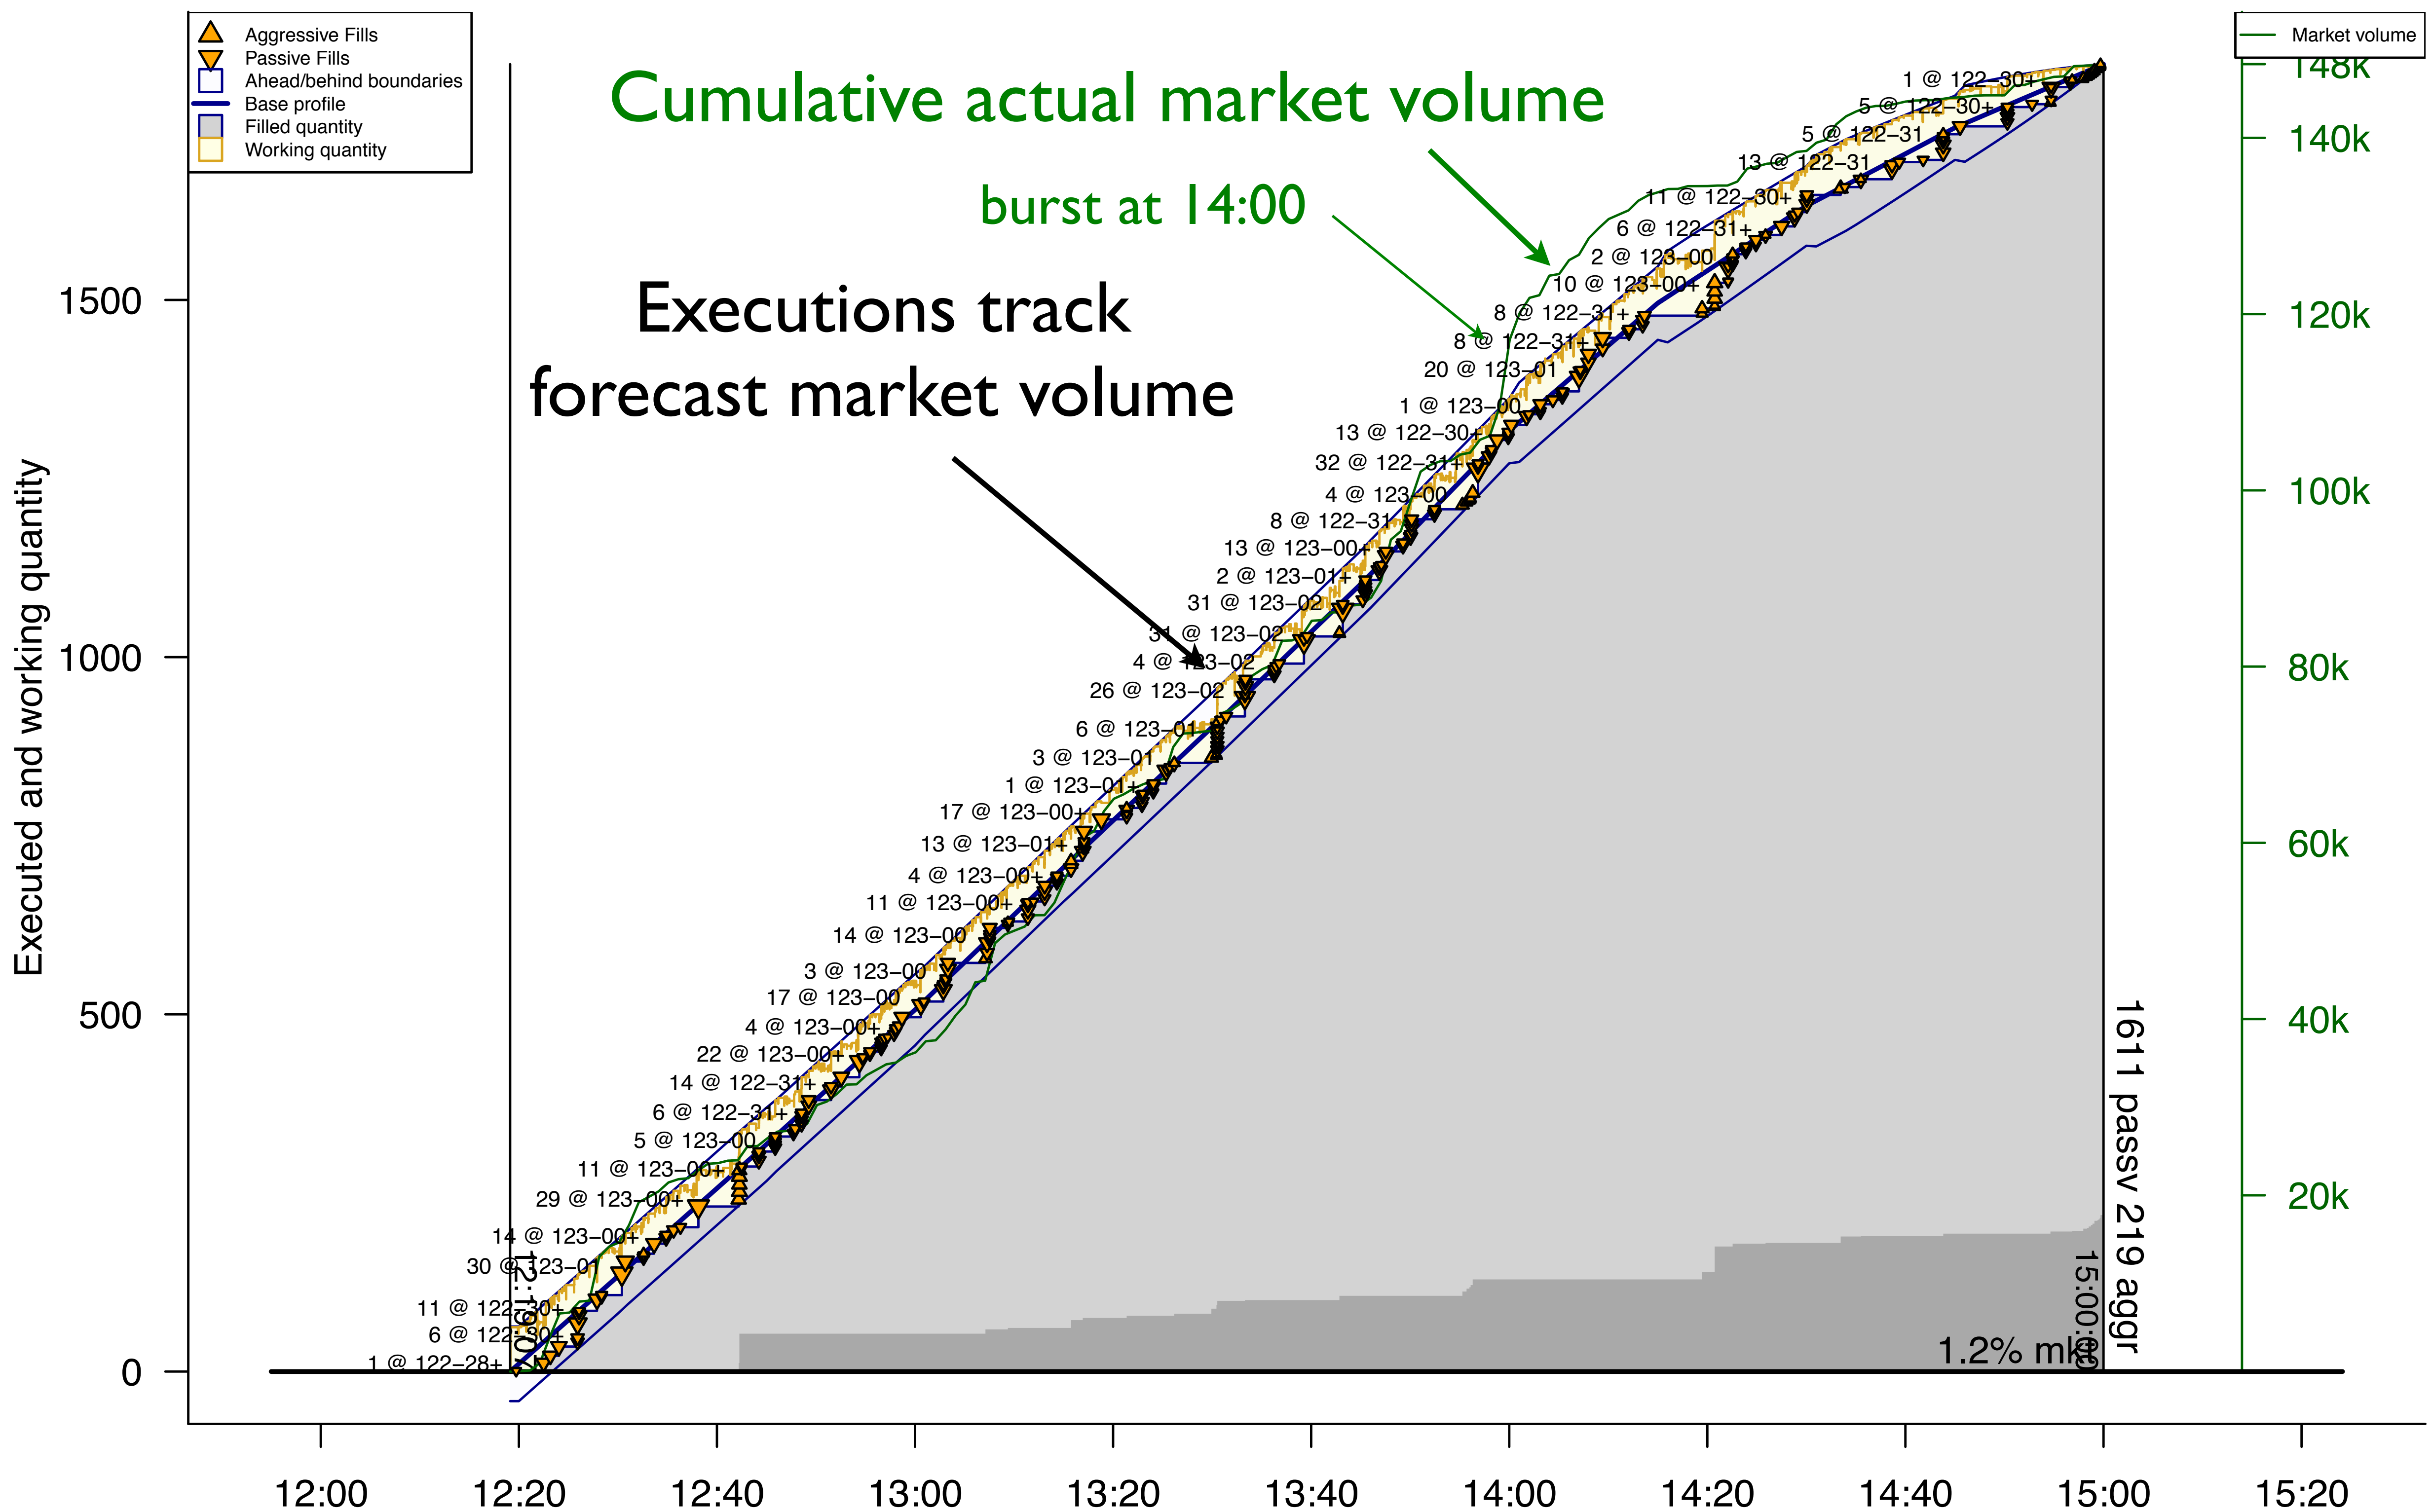

Produced by Quantitative Brokers from internal and CME Market Data

VWAP Example (2)

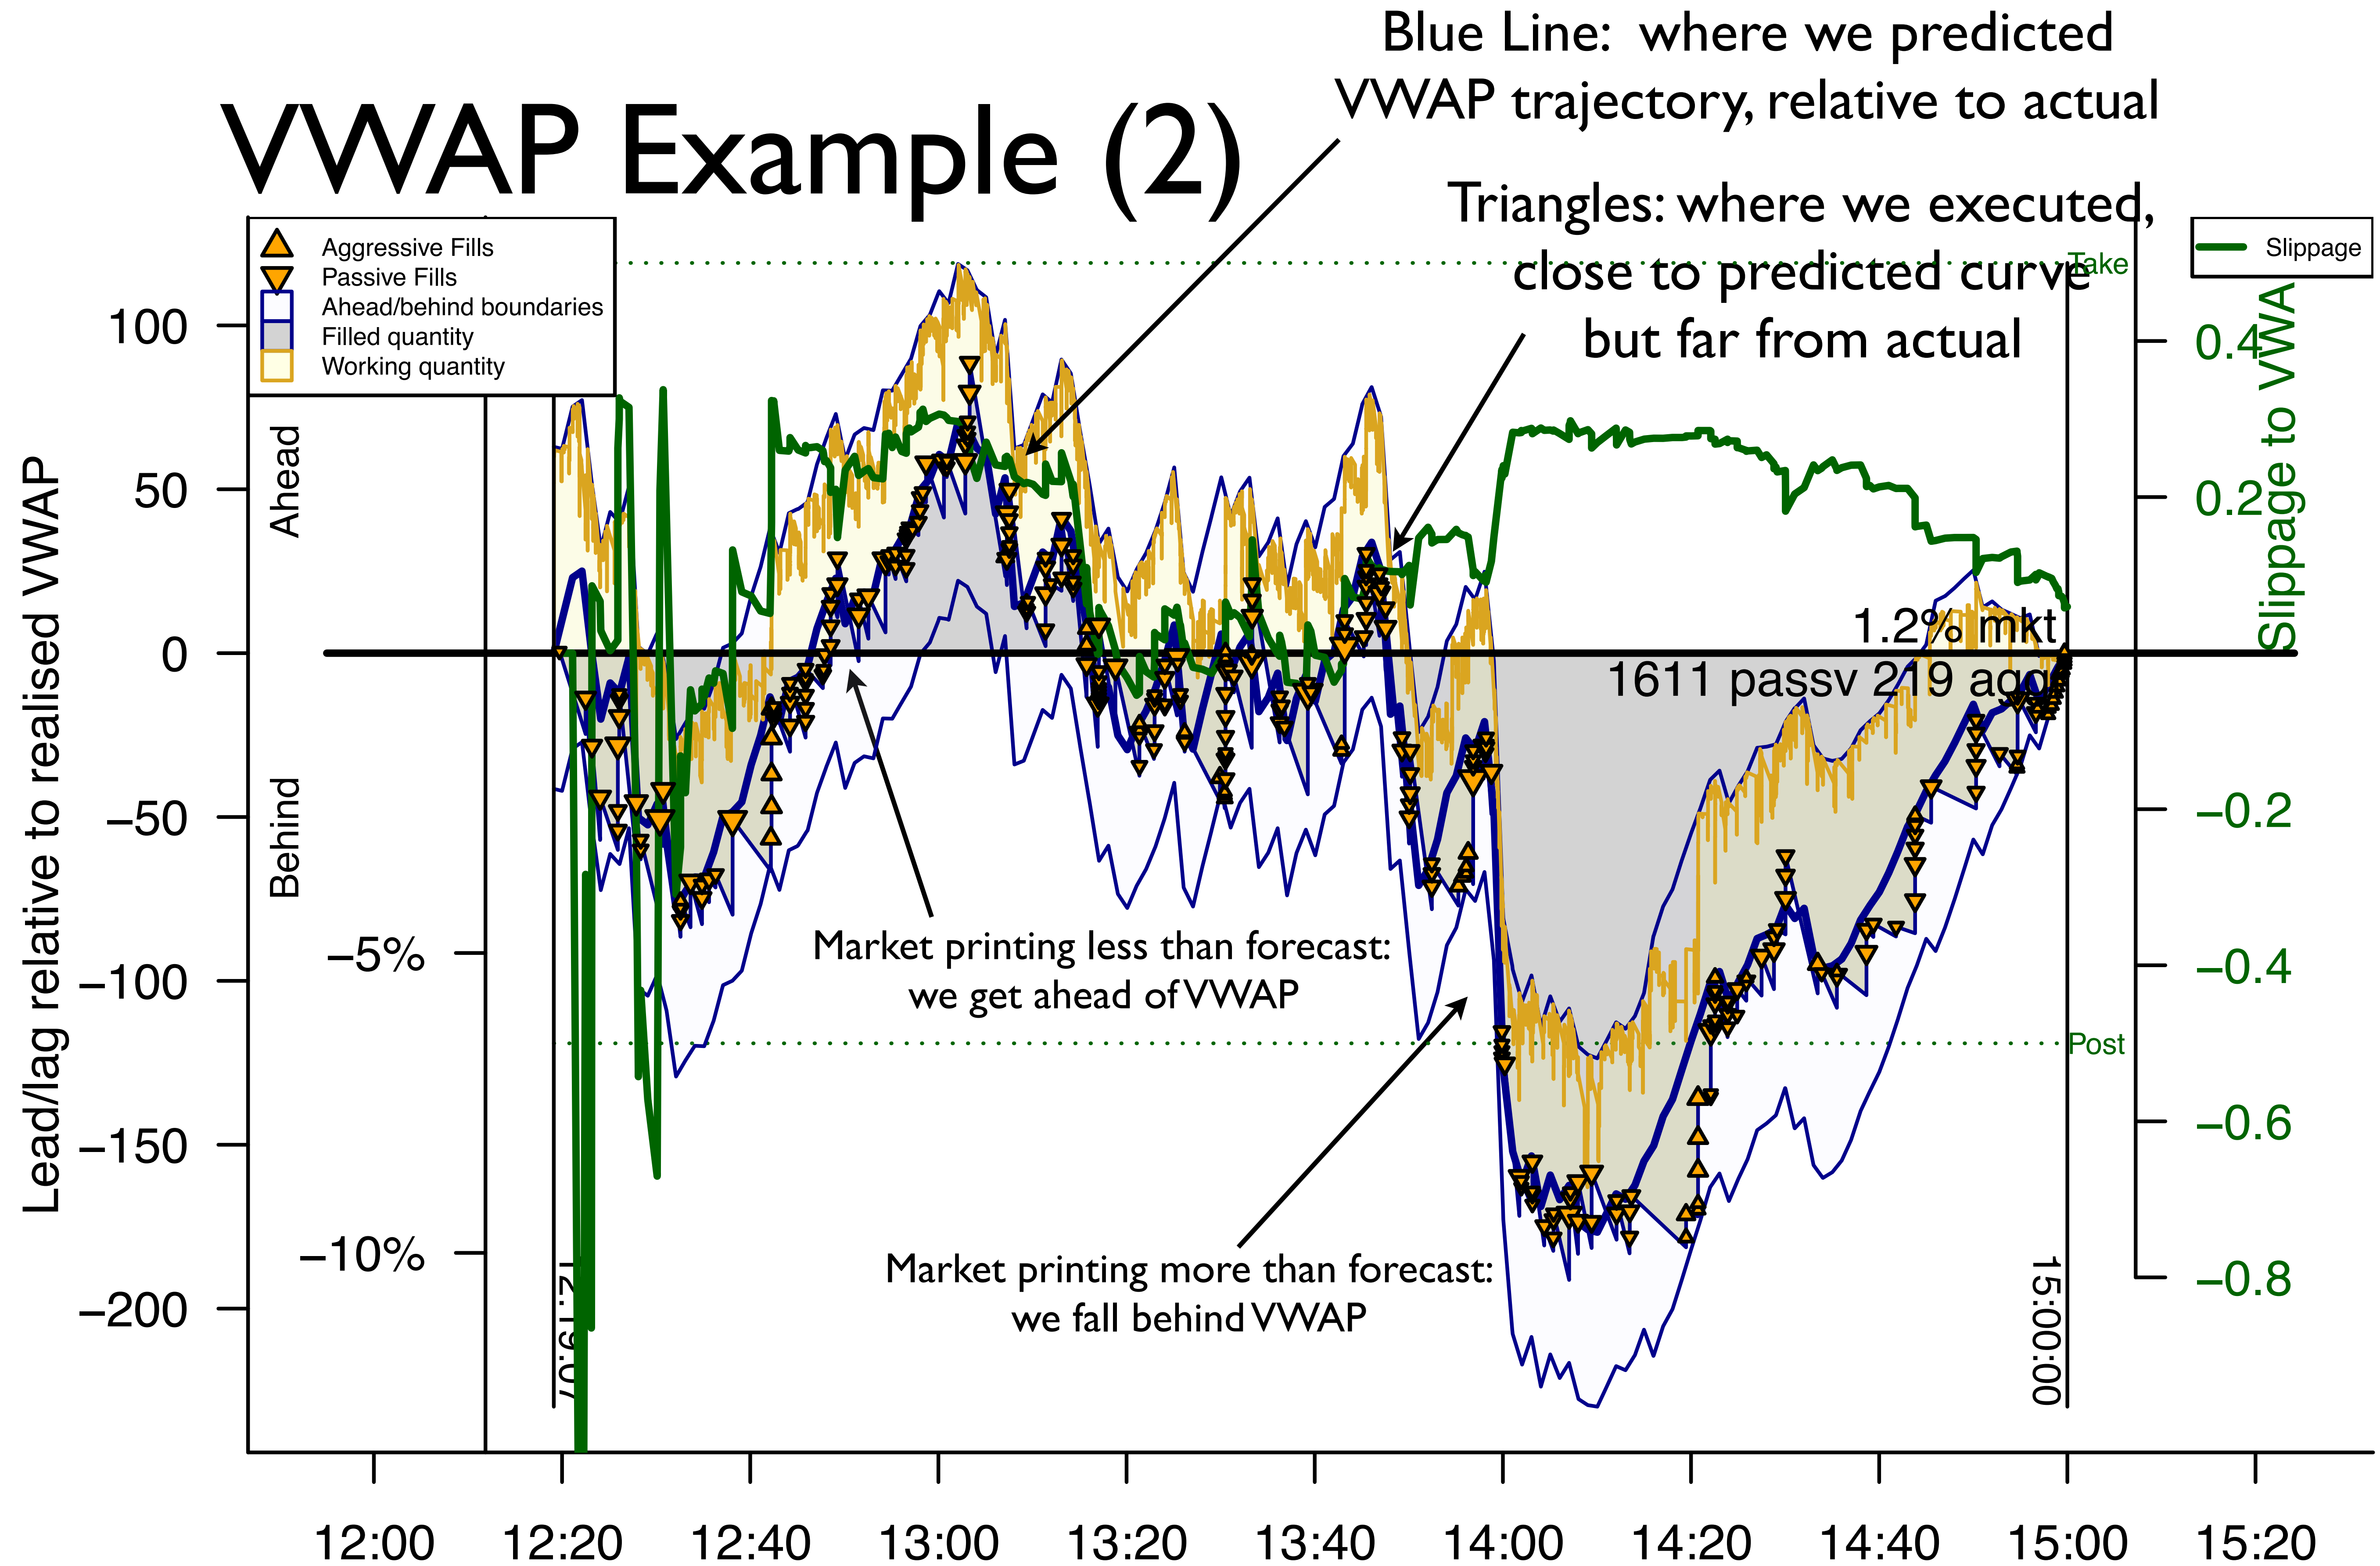

TWAP benchmark

Time-Weighted Average Price

- ignoring trade volume

Positives (relative to VWAP):

- easy to achieve
- easy to measure where volume info is unreliable

Negatives (relative to VWAP):

- not related to market activity

Minimize slippage to TWAP benchmark
by trading exactly along straight-line profile

TWAP Example (1)

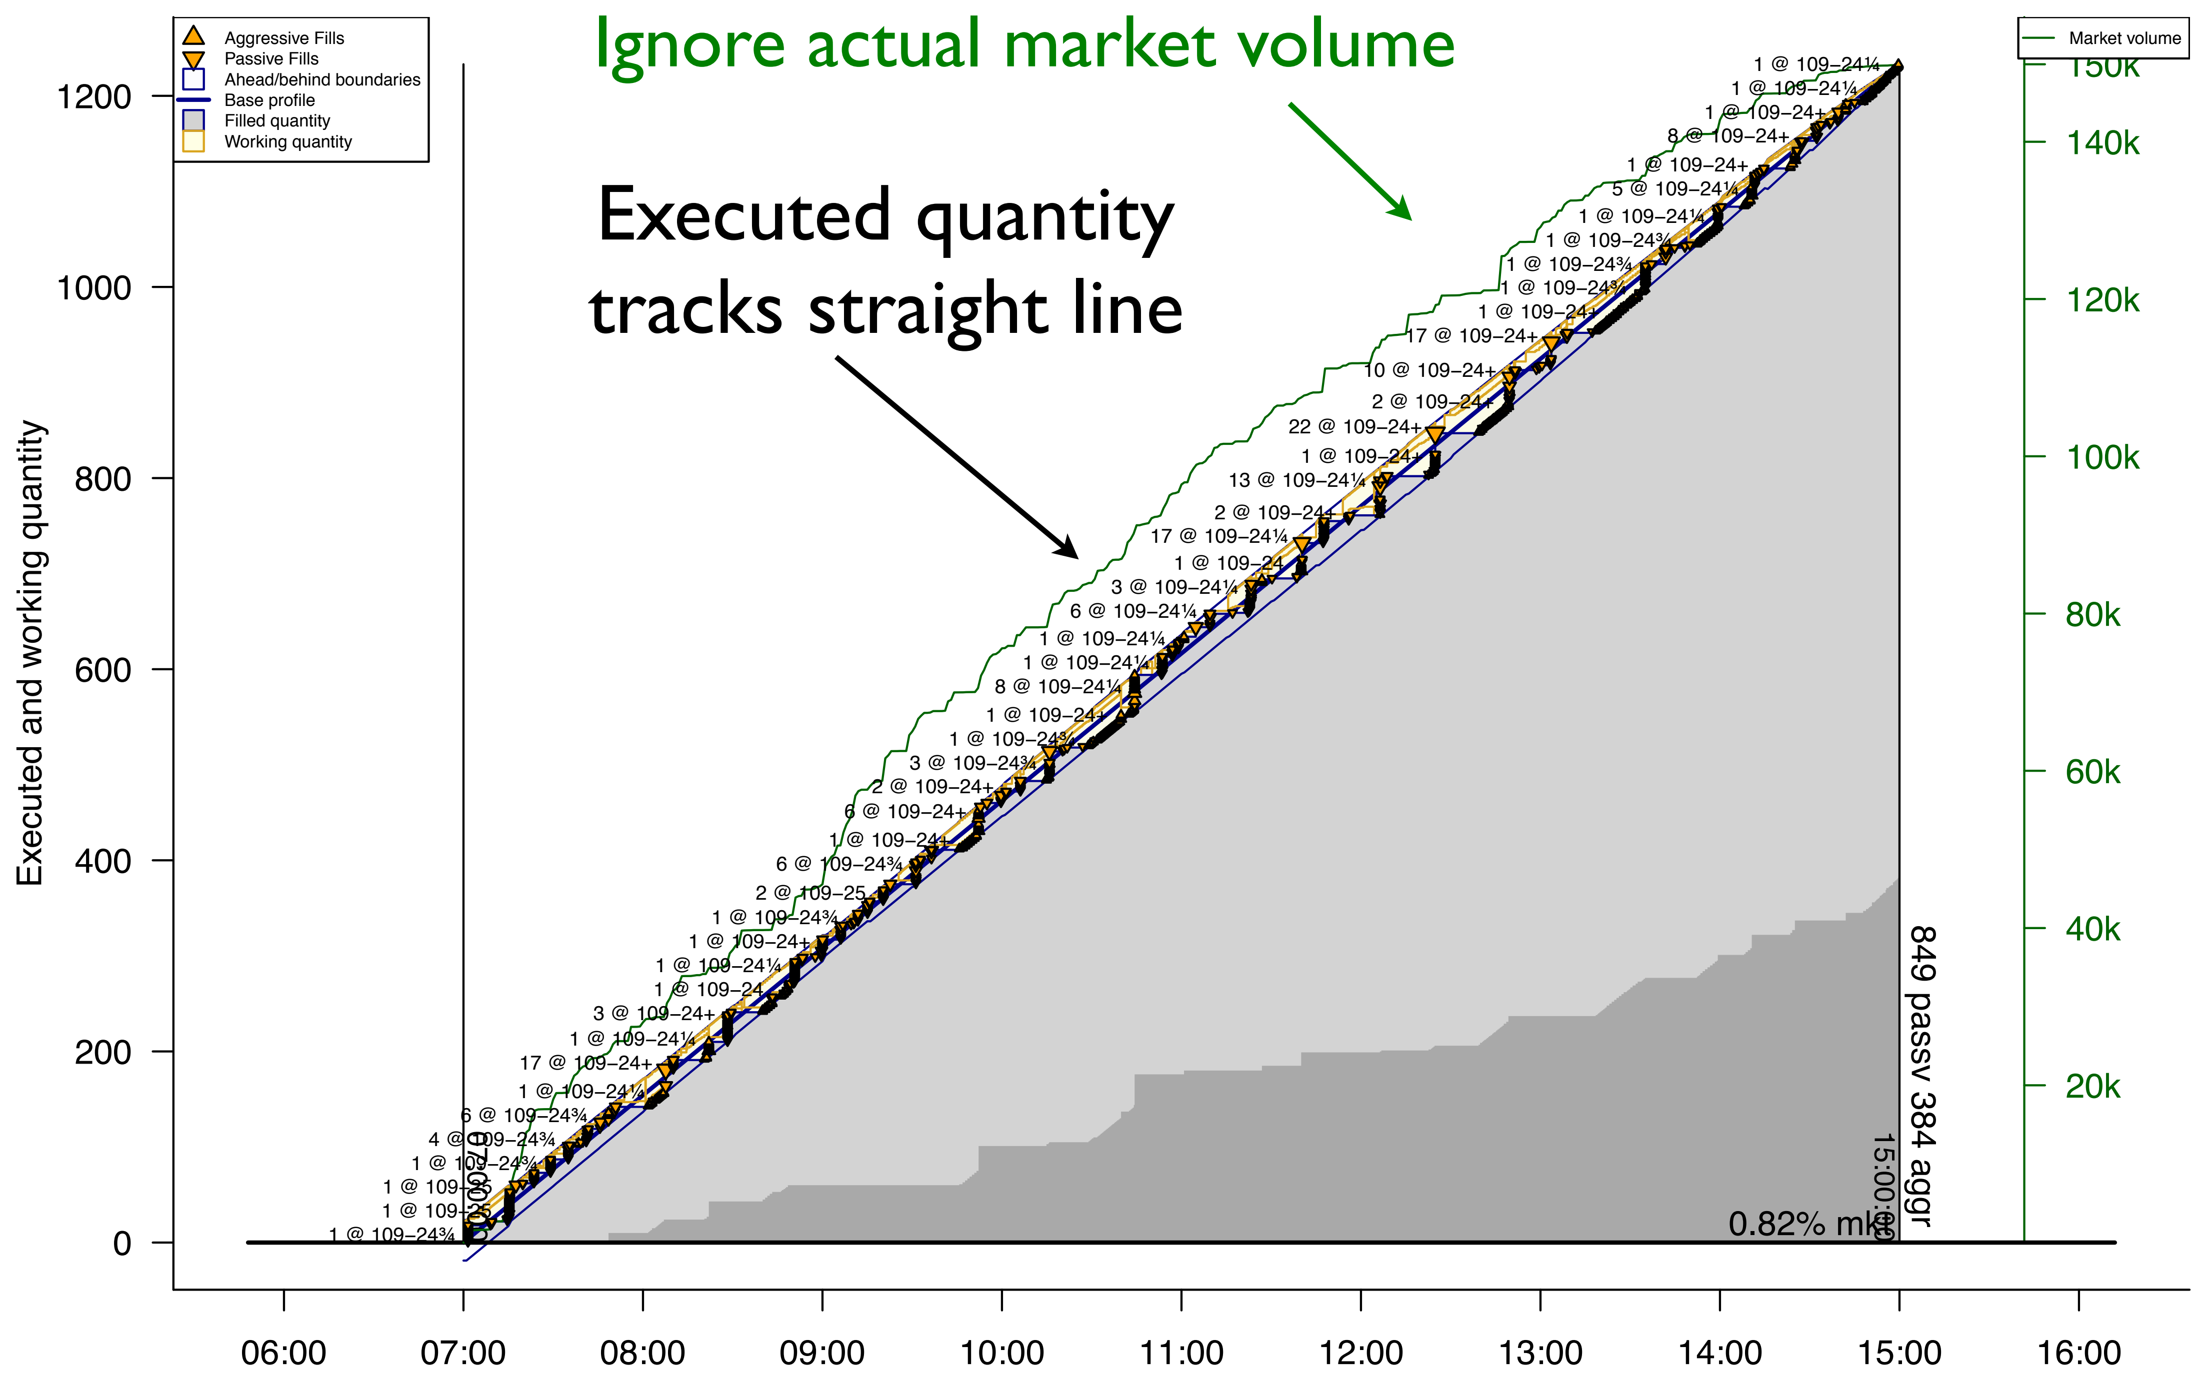

Produced by Quantitative Brokers from internal and CME Market Data

TWAP example (2)

Cumulative fill quantity consistently within ±1% of true TWAP

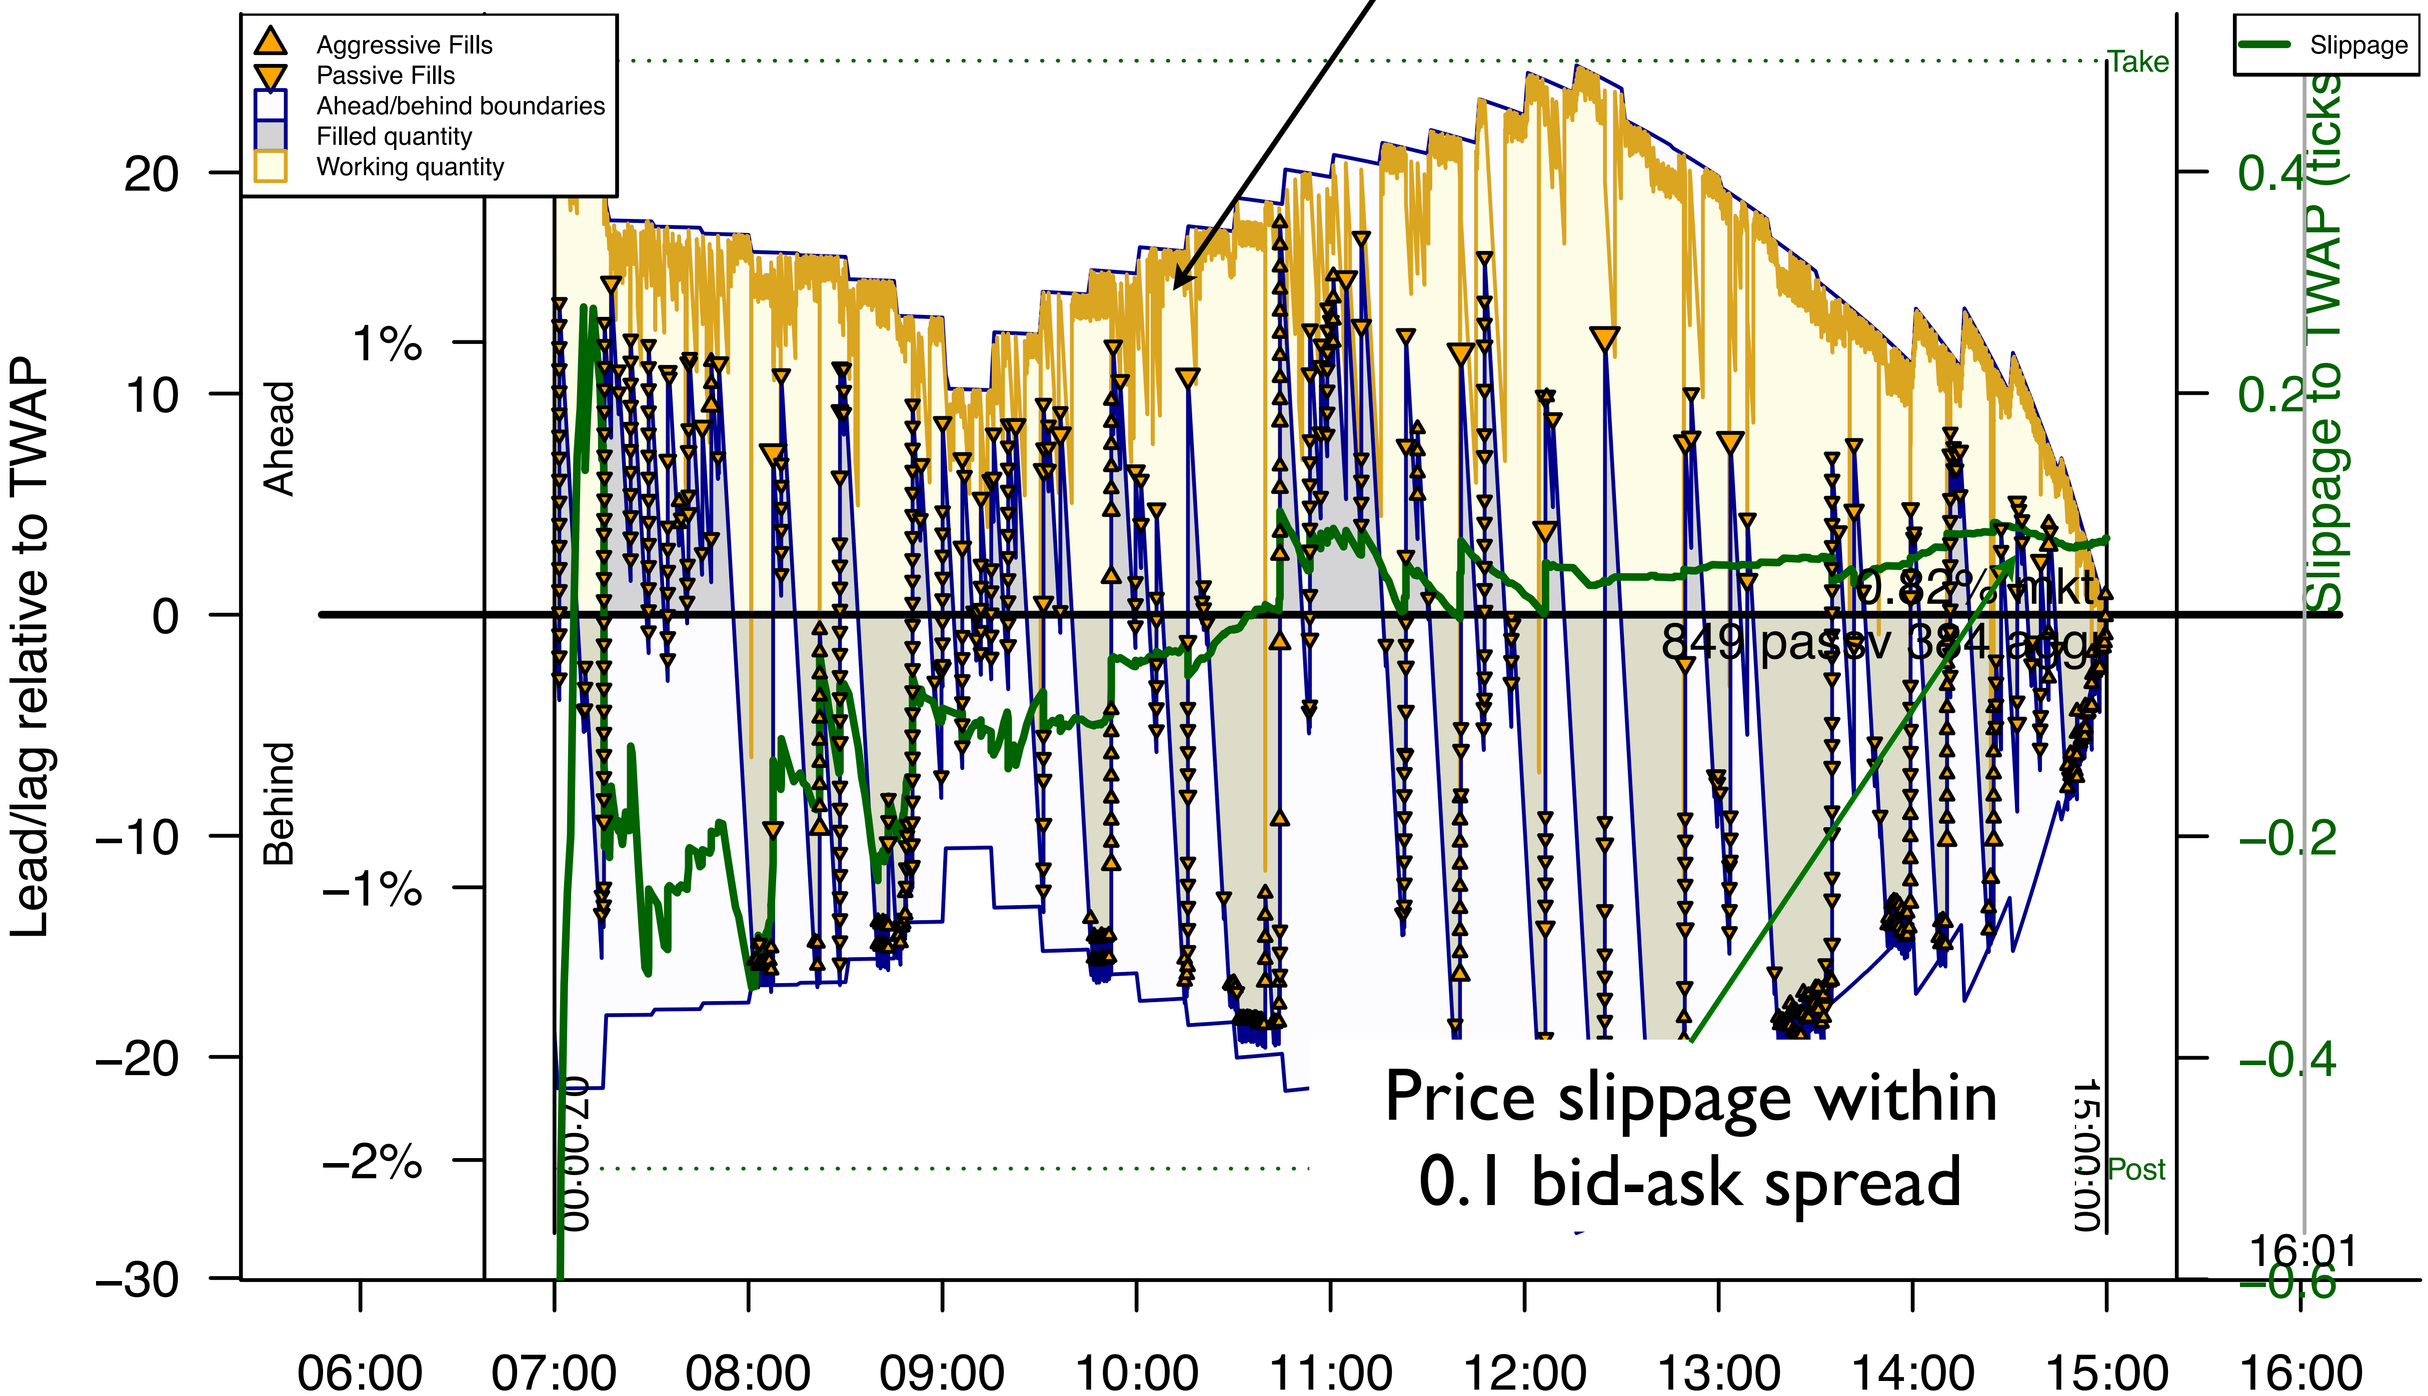

Price slippage within 0.1 bid-ask spread

Produced by Quantitative Brokers from internal and CME Market Data

Other benchmarks

Daily close price, or settlement

  without impacting it too much

Option hedging

  achieve delta-neutral position, depending on price
  (work with Tianhui Li, Princeton)

Etc.

Benchmark must be agreed in advance
Execution is tailored to benchmark
Results must be evaluated against same benchmark

Execution is uncertain

Market price moves during execution
Limit orders may or may not get filled
Liquidity appears and disappears
Volume profiles are unpredictable

Result of execution is a “random variable”
Execution strategy tailors its properties
e.g. mean and variance

Perold's "cost of not trading"

Fundamental Balance of Execution

Cost ⟷ vs ⟷ **Risk**

Cost		Risk
(mean of executed price relative to benchmark)	vs	(variance of executed price relative to benchmark)
Alpha (forecastable price change) Passive fill probability Market impact		Price volatility (most important)

Example #1: Arrival Price

Minimize risk: execute as quickly as possible
But
Too fast = lower chance of passive fills
  when to cross spread to reduce risk?
Too fast = market impact pushes price
  how much is it worth it to push price?
Alpha signal says price will get worse
  how much should you accelerate? increase impact
Alpha signal says price will improve
  how much slower should you trade? increase risk

Example #2: TWAP

Ideal profile: evenly spaced in time

Should you cross spread whenever needed
or wait for passive fills even if behind?

Alpha signal: how much should you deviate from
schedule to exploit signal?
(Signal may be very uncertain)

Example #3: VWAP

Track volume forecast, fixed or dynamic

Should you worry less about deviating from schedule, since schedule itself is uncertain?

2. Price dynamics and market impact

1. Mathematical model for trading and impact
2. Definition of objective function
3. Mathematical optimization

Trading trajectory

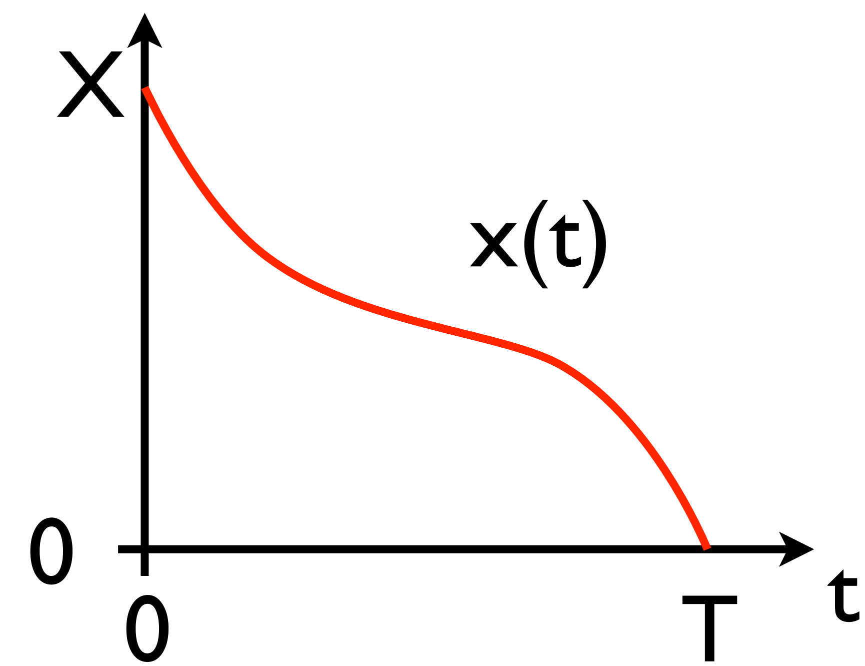

X = target number of shares to buy
T = time limit
$x(t)$ = shares remaining at time t

$$x(0) = X, \; x(T) = 0$$

$v(t) = -dx/dt$ = rate of buying
$x(t), v(t)$ can depend on all information to time t
or be fixed in advance (adaptive / non-adaptive)
Detailed orders are placed by micro-algorithm

Market model (Almgren/Chriss '00)

$S(t)$ = stock price

$$dS(t) = \sigma(t)\, dB(t) \;+\; g(v(t))\, dt$$

$\sigma(t)$ = volatility

$$g(v) \;=\; \text{permanent market impact}$$

Arbitrage exists unless g(v) is linear

$$g(v) \;=\; \gamma v$$

Integrates out so neglect permanent impact

Temporary market impact

$$\tilde{S}(t) \; = \; \text{realized trade price}$$

$$\tilde{S}(t) \; = \; S(t) \; + \; h(v(t))$$

$$\text{Linear case:} \quad h(v) \; = \; \eta(t)\, v$$

Market parameters σ(t) and η(t)

- constant, predictable, or random processes
- observable in real time

Temporary market impact

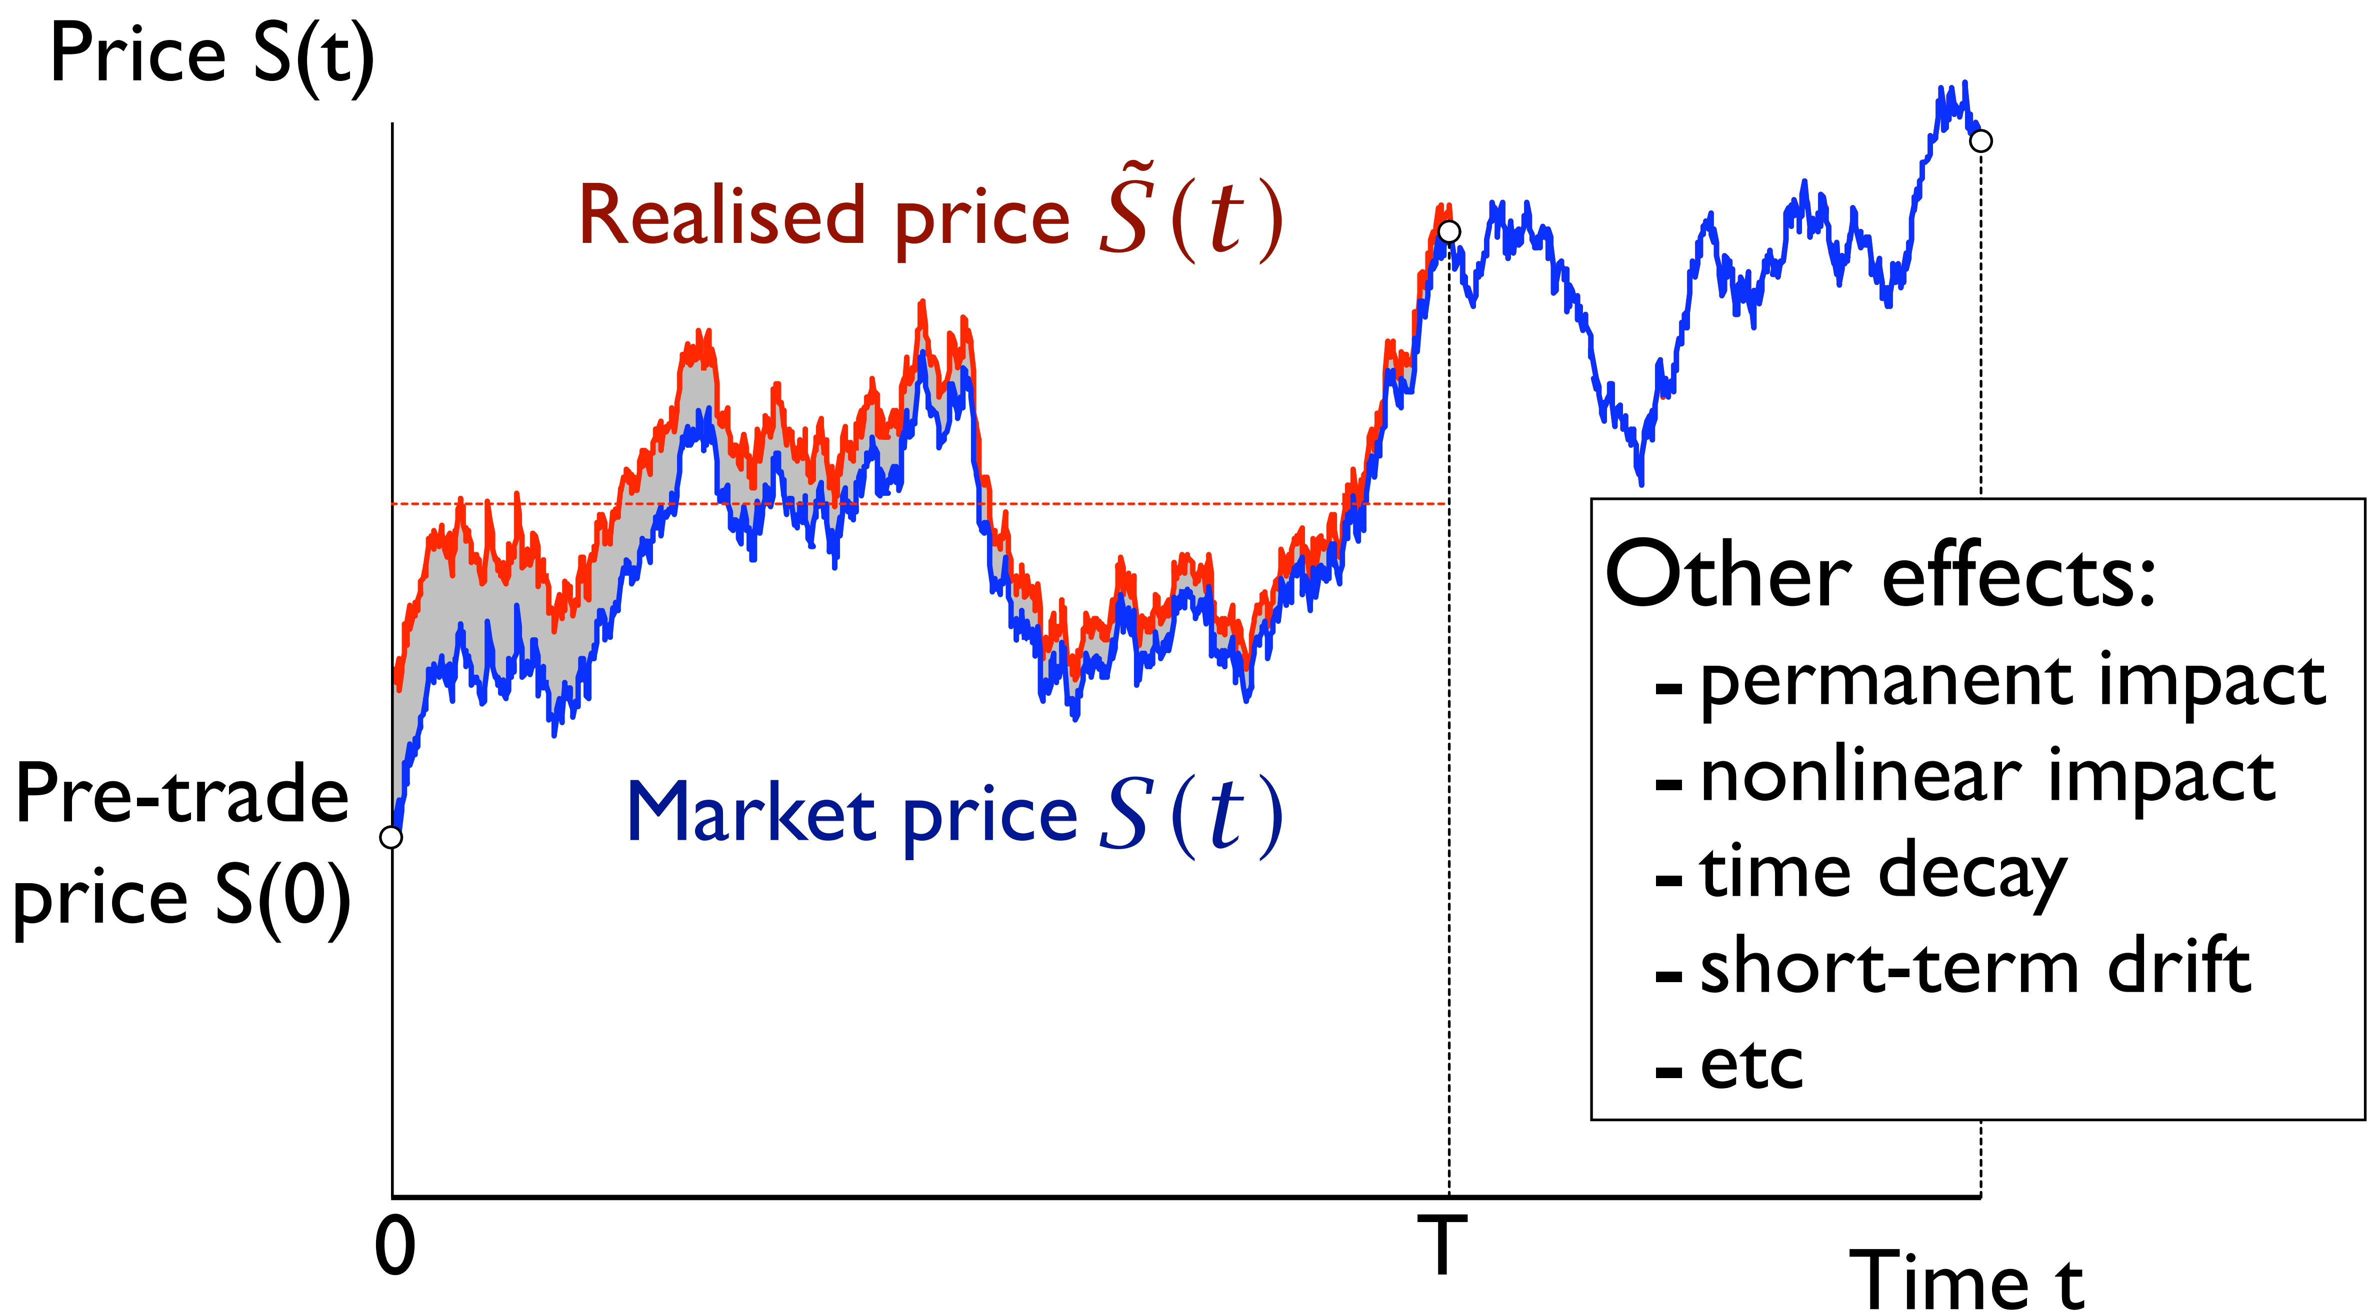

Cost of trading

"Capture" = Dollars realized on selling X shares

$$= \int_0^T \tilde{S}(t)\, v(t)\, dt$$

$$= X S_0 + \int_0^T \sigma(s) x(s)\, dB(s) + \int_0^T \eta(t)\, v(t)^2\, dt$$

Cost = capture - initial value X S_0

$$C = \int_0^T \sigma(s) x(s)\, dB(s) + \int_0^T \eta(t)\, v(t)^2\, dt$$

Random because of price change and trade trajectory decisions

3. Mean variance optimization

Minimize both

E = expectation of C

V = variance of C

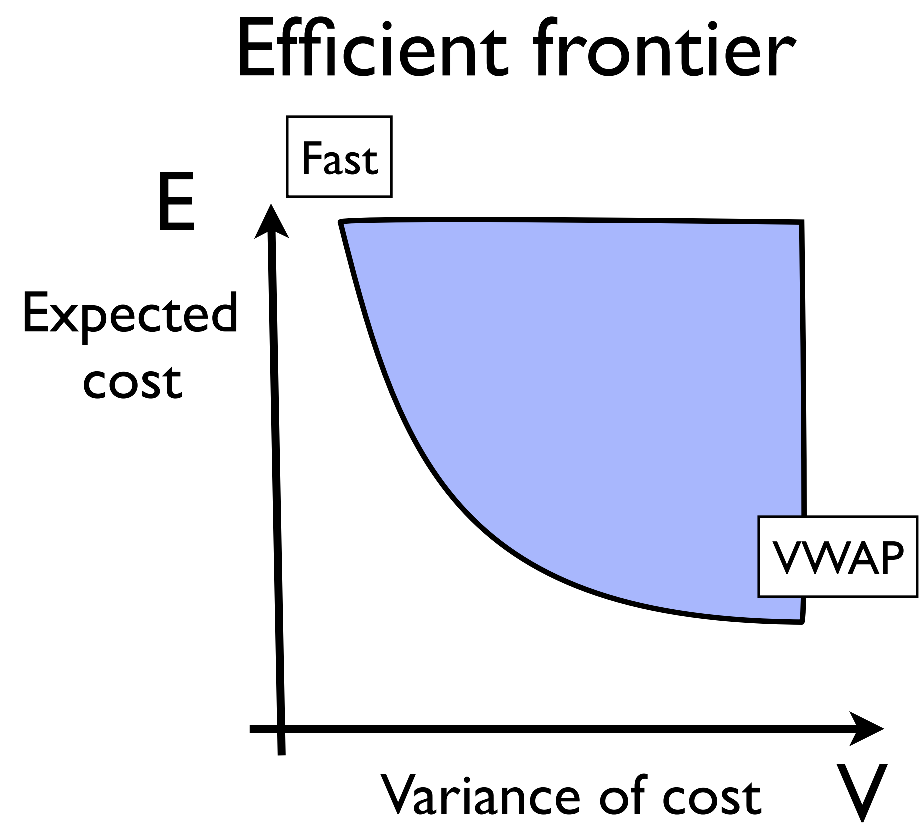

At each time t

$$\min_{x(s):t\le s\le T} \mathbb{E}(C) + \lambda \, \mathrm{Var}(C)$$

Risk parameter λ set by client

Solution

$$C = \sigma \int_0^T x(t)\, dB(t) + \eta \int_0^T v(t)^2\, dt$$

$$V = \sigma^2 \int_0^T x(t)^2\, dt, \qquad E = \eta \int_0^T v(t)^2\, dt$$

(for a precomputed fixed trajectory, which is optimal solution)

Calculus of variations with quadratic objective:

Exponential solution

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

$$v(s) = \kappa\, x(t)\, \frac{\cosh(\kappa(T-s))}{\sinh(\kappa(T-t))}$$

Time scale

$$\kappa^2 = \frac{\lambda\sigma^2}{\eta}$$

Arrival price solution

Time constant

$$\kappa^2 = \frac{\lambda\sigma^2}{\eta}$$

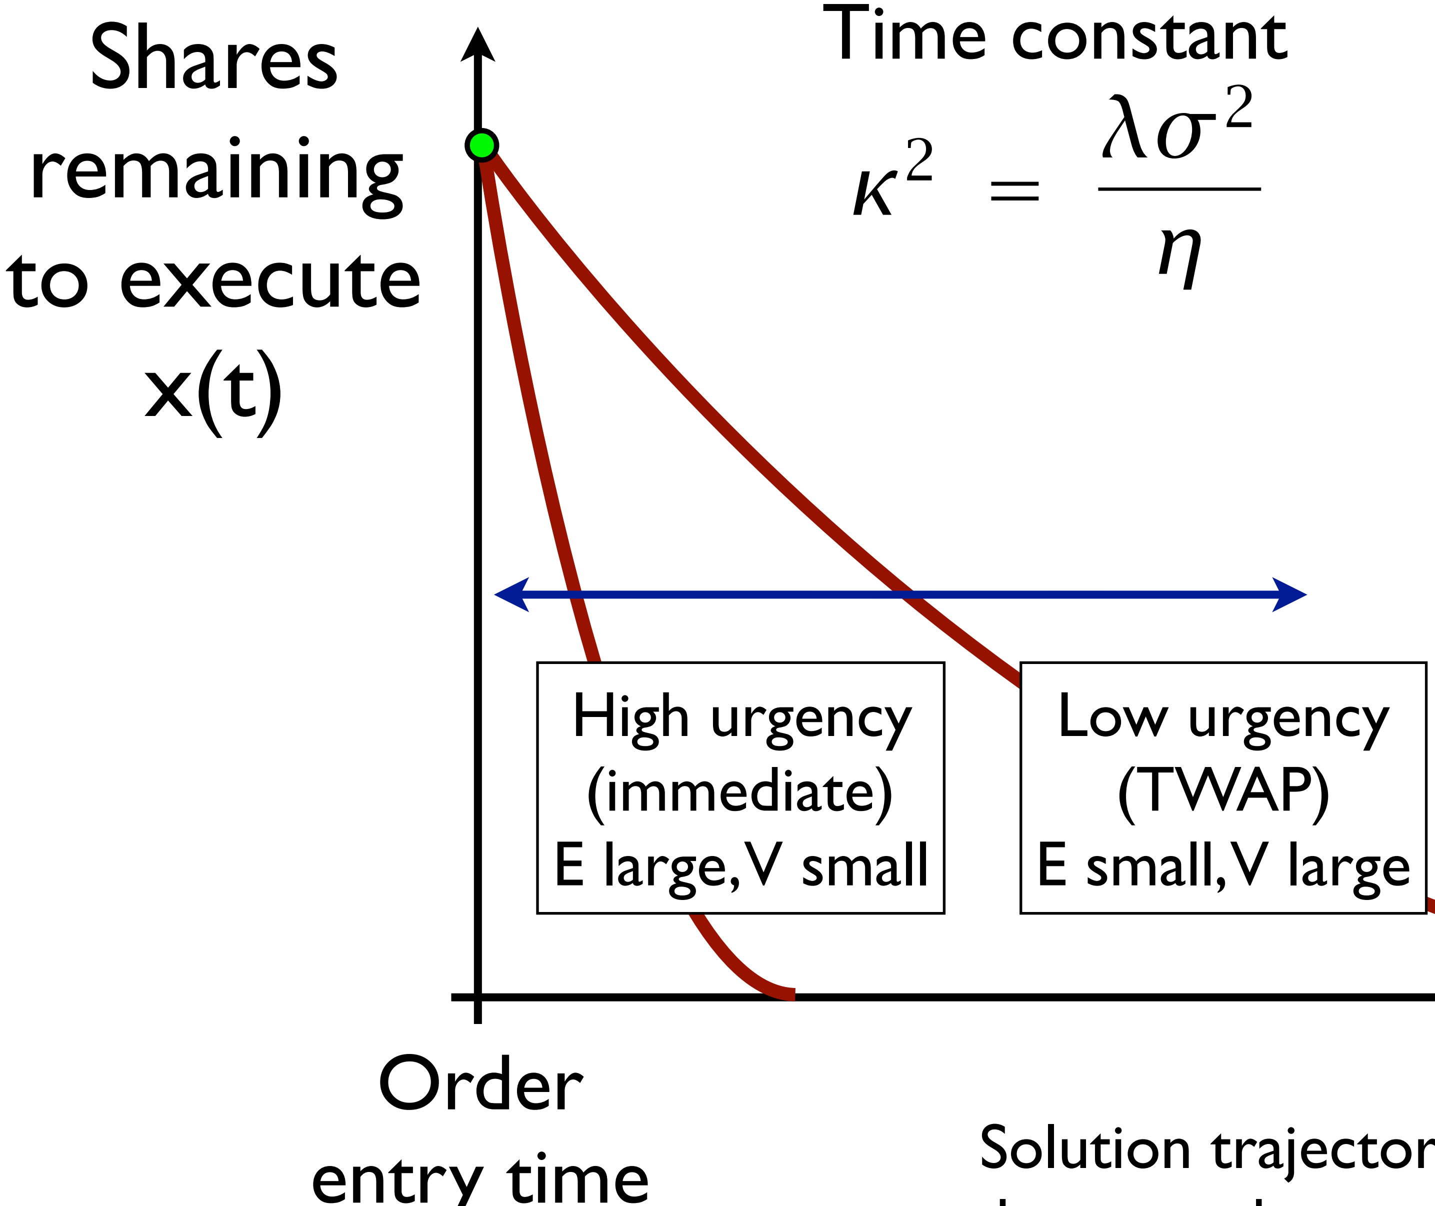

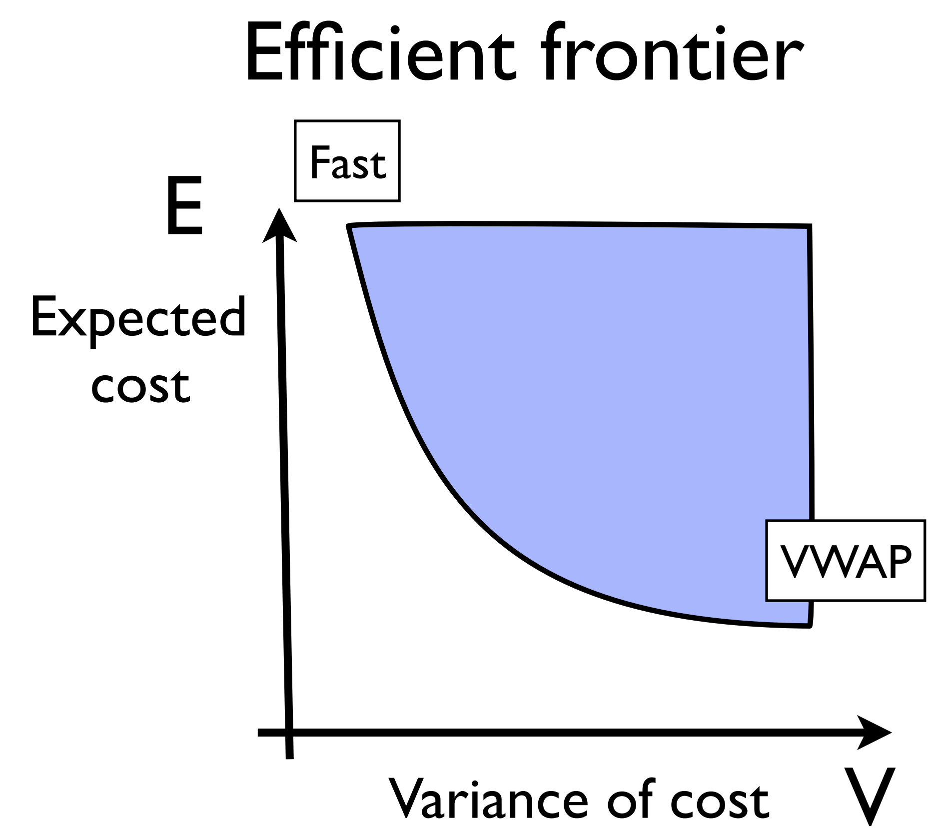

Solution trajectory does not change as price motions are observed during execution

Extensions

Portfolios (Almgren/Chriss 2000)
Nonlinear market impact (Almgren 2003)
Adaptive trajectories (Lorenz/Almgren 2011)
Stochastic liquidity and volatility (Almgren 2012)
Bayesian update price drift (Almgren/Lorenz 2006)
Option hedging (Almgren/Li 2016)
Many others ...

Critique of arrival price solutions

Well calibrated market impact model

Known & constant risk aversion parameter λ

Reasons to trade slowly: reduce market impact

- slow steady trading leaks information
- price drift can lead to trade faster
- opportunistic trading

Reasons to trade fast: reduce volatility risk

- event risk
- reduce number of open orders

Market impact is more complicated than this

- persistence in time (integral kernels)
- bursty trading can be cheaper than steady

4. Option hedging

Black-Scholes with market impact
Effect on public markets

Market Microstructure and Liquidity
Vol. 2, No. 1 (2016) 1650002 (26 pages)

DOI: 10.1142/S2382626616500027

Option Hedging with Smooth Market Impact

Robert Almgren*
Quantitative Brokers, New York, NY, USA
ralmgren@quantitativebrokers.com

Tianhui Michael Li
The Data Incubator, New York, NY, USA

Equity price swings on July 19 2012

(one day prior to options expiration)

Bloomberg Our Company | Professional | Anywhere

HOME QUICK NEWS OPINION MARKET DATA PERSONAL FINANCE TECH POLITICS

http://www.bloomberg.com/news/2012-07-20/hourly-price-swings-whipsaw-investors-in-ibm-coke-mcdonald-s.html

Investors Whipsawed by Price Swings in IBM, Coca-Cola

By Inyoung Hwang, Whitney Kisling & Nina Mehta - Jul 20, 2012 5:22 PM ET

Investors in three of the biggest Dow Jones Industrial Average (INDU) stocks were whipsawed by price swings that repeated every hour yesterday, fueling speculation the moves were a consequence of computerized trading.

Shares of International Business Machines Corp. (IBM), McDonald's Corp. (MCD) and Coca-Cola Co. (KO) swung between successive lows and highs in intervals that began near the top and bottom of each hour, data compiled by Bloomberg show. While only IBM finished more than 1 percent higher, the intraday patterns weren't accompanied by any breaking news in the three companies where $3.42 billion worth of shares changed hands.

Marko Kolanovic, global head of derivatives and quantitative strategy at JPMorgan & Chase Co

• • •

The four stocks with repeating price patterns yesterday had the biggest net long options positions among S&P 500 Index companies, according to JPMorgan's calculations, Kolanovic said in a note to clients today. The amount traded in the stocks was also consistent with what traders would have had to buy or sell, indicating that the patterns could be "almost entirely explained" by their need to hedge, he wrote.

"We believe that the price pattern of KO, IBM, MCD and AAPL yesterday was caused by hedging of options by a computer algorithm," Kolanovic said in the note that referred to the companies by their ticker symbols. "It was likely an experiment in automatically hedging large option positions with a time- weighted algorithm that has gone wrong for the hedger."

CHEUVREUX
CRÉDIT AGRICOLE GROUP

QUANT **NOTE**

28th of August, 2012

What does the saw-tooth pattern on US markets on 19 July 2012 tell us about the price formation process

The saw-tooth patterns observed on four US securities on 19 July provide us with an opportunity to comment on common beliefs regarding the market impact of large trades; its usual smoothness and amplitude, the subsequent "reversal" phase, and the generic nature of market impact models.
This underscores the importance of taking into account the motivation behind a large trade in order to optimise it properly, as we already emphasised in *Navigating Liquidity 6*.

We used different intraday analytics to work out what happened: pattern-matching techniques, market impact models, order flow imbalances and PnL computations of potential stat. arb intraday strategies. After looking at open interests of derivatives on these stocks, we conclude that repetitive automated hedging of large-exposure derivatives lay behind this behaviour. This is an opportunity to understand how a very crude trading algorithm can impact the price formation process ten times more than is usually the case.

Charles-Albert LEHALLE
Global Head of Quantitative Research
clehalle@cheuvreux.com
(33) 1 41 89 80 16

Matthieu LASNIER
Market microstructure specialist, Quantitative Research, Paris

Paul BESSON
Index and Portfolio Research Specialist, Quantitative Research, Paris

Hamza HARTI
Trading specialist, *Quantitative Research, Paris*

Weibing HUANG
PhD student, *Quantitative Research, Paris*

Nicolas JOSEPH
Performance analysis specialist, *Quantitative Research, London*

Lionel MASSOULARD
Trading specialist, *Quantitative Research, New-York*

Chevreux: sawtooths caused by options hedging

FIGURE 18: SAW-TOOTH ON COCA-COLA AND 77.50 CALL

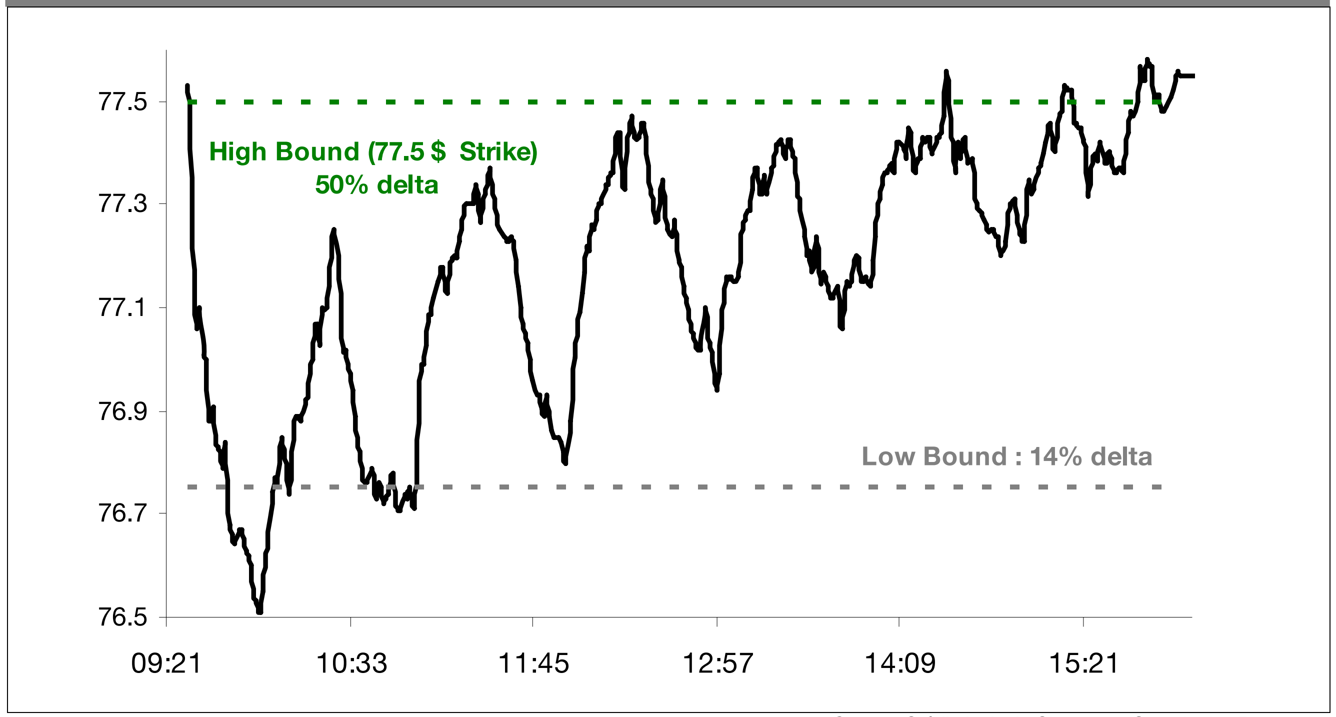

Source: Crédit Agricole Cheuvreux Quantitative Research

FIGURE 19: LONG VOLATILITY DELTA HEDGING

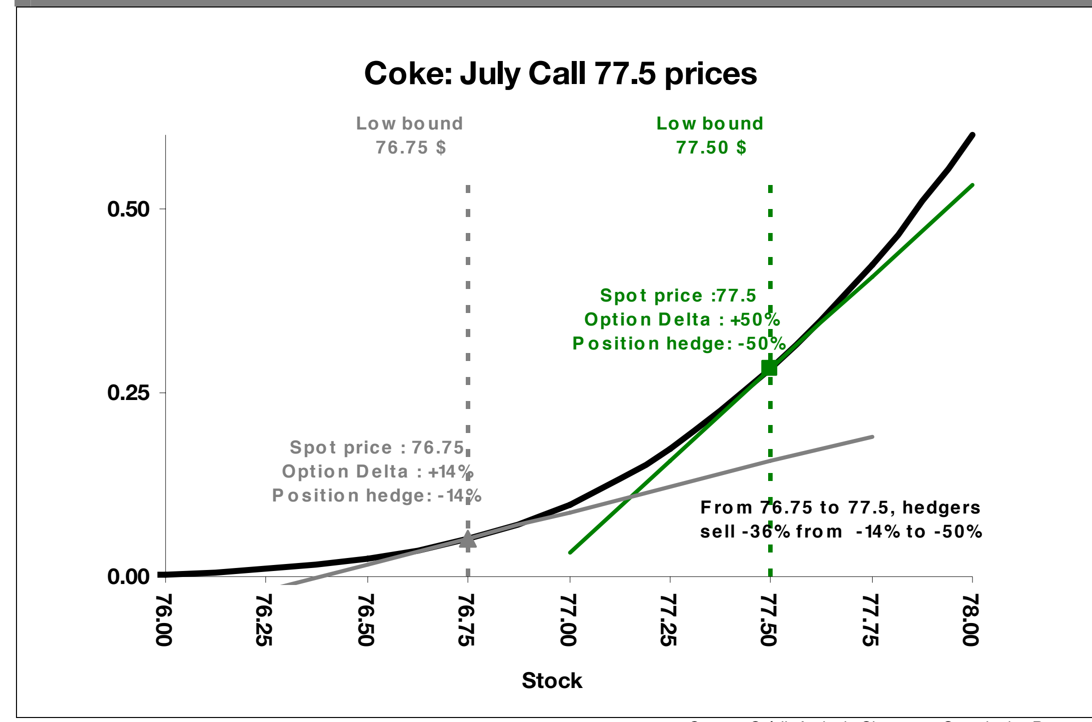

Source: Crédit Agricole Cheuvreux Quantitative Research

For a very large open interest position encompassing a large gamma, a significant move in the stock price will have disastrous effects for a basic rudimental hedger such as the one described above.

- **Repetitive delta hedging seems to be the most plausible explanation**

Simplistic hedging of large gamma options is a plausible explanation for the "saw-tooth" trading pattern. This explanation is consistent with the main features of this phenomenon: timing, aggressiveness, impact, predictability and information leakage which is what characterises those "saw-tooth" patterns. Fortunately, large option positions are most often managed dynamically in a continuous way, and discrete archaic hedging processes have almost disappeared in modern-day markets.

(we show how to do this better)

Our solutions with proportional cost

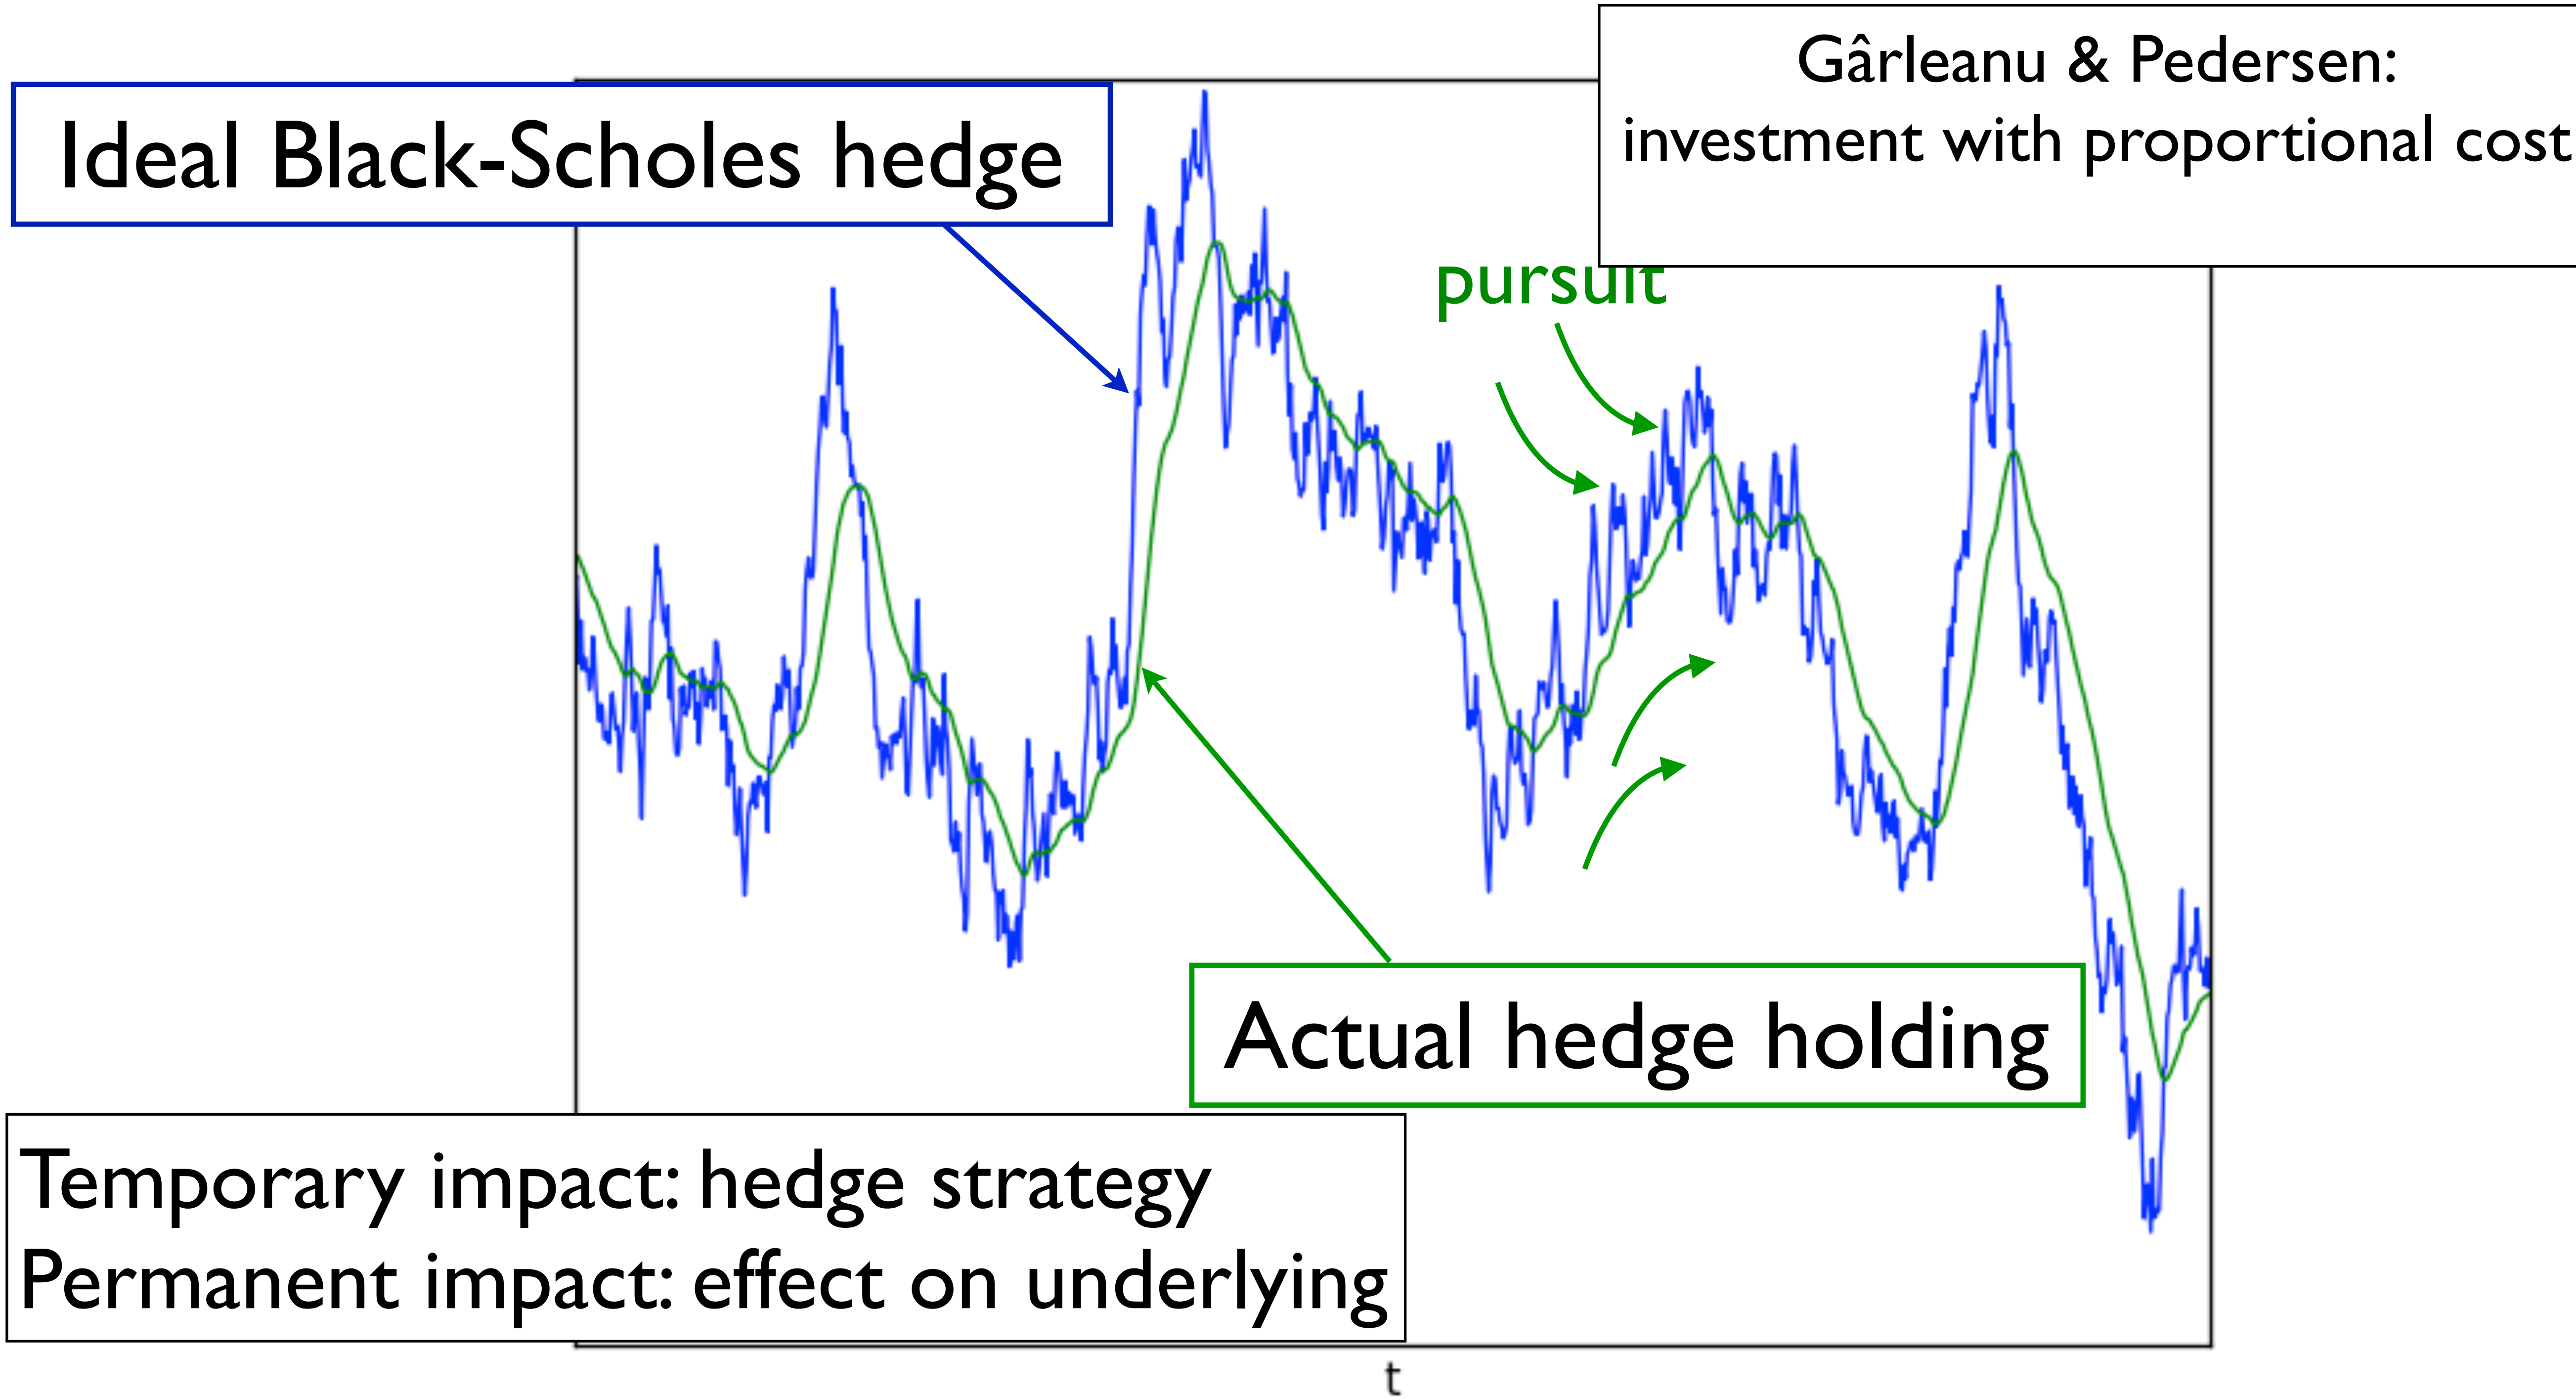

$$\theta_t \;=\; -\kappa\, h(\kappa(T-t)) \cdot (X_t - \text{target})$$

What happens to price process?

$$Z_t = -\nu(X_t - X_0) + (P_t - P_0)$$

$$dZ_t = \sigma\, dW_t \longleftarrow \text{same as in } P_t$$

$$P_t = P_0 + \frac{1}{1+\nu\Gamma}(\nu Y_t + Z_t)$$

$$= P_0 + \frac{\sigma}{1+\nu\Gamma}\left(W_t + \nu\Gamma \int_0^t e^{-\kappa(1+\nu\Gamma)(t-s)} dW_s \right)$$

modified volatility

stationary $\propto 1/\sqrt{\kappa}$

Γ<0: hedger is short the option

1+νΓ<1: overreaction, increased volatility

Γ>0: hedger is long the option

1+νΓ>1: underreaction, reduced volatility

“Signature plot”

$$\tilde{\kappa} = (1 + \nu\Gamma)\kappa$$

unconditional (F_0)

$$\frac{1}{s-t}\mathbb{E}(P_s - P_t)^2 =$$

$$= \left(\frac{\sigma}{1+\nu\Gamma}\right)^2 \left(1 + (2+\nu\Gamma)\nu\Gamma\,\frac{1-e^{-\tilde{\kappa}(s-t)}}{\tilde{\kappa}(s-t)}\right)$$

Volatility measured on time interval Δt

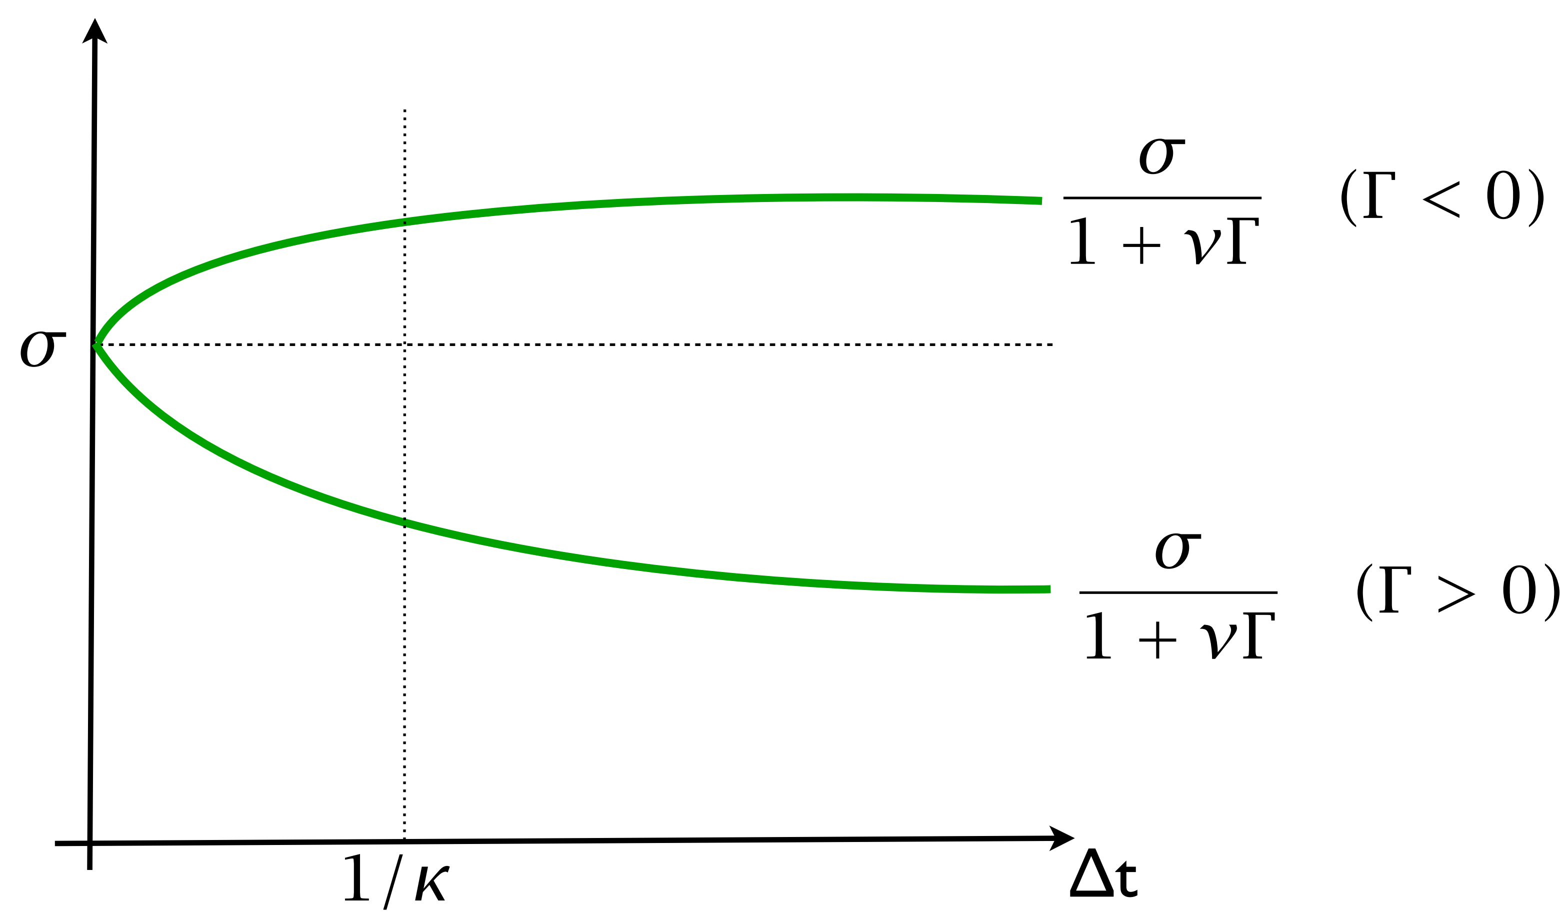

Nondimensional parameter

$$1 + \nu\Gamma$$

$$\nu = \frac{\text{price change due to market impact}}{\text{shares executed}}$$

$$\Gamma = \frac{\text{shares needed to adjust hedge}}{\text{market price change}}$$

Problems unless $|\nu\Gamma| \lesssim 1$

Conclusions

Many interesting math models

- calculus of variations
- dynamic programming

Models are very approximate

- markets are messy

Can give insight nonetheless

- trade fast vs trade slow
- permanent and temporary impact