Market Microstructure and Algorithmic Trading

PIMS Summer School 2016
University of Alberta, Edmonton
Lecture 3: July 5, 2016

Robert Almgren
quantitativebrokers

Outline

1. Micro impact
   individual child orders
   paid by liquidity demander to supplier
2. Macro impact
   need parent orders (brokers only)
   incorporate time or no time
3. Models for trade optimization

Market impact

How your trades affect the market

How much it costs you to trade

"micro" impact: individual trades or events

- execute trade with market order
- or place/cancel limit order

"macro" impact: larger scale orders

- "buy 1000 lots across next 2 hours"
- how does price change during and after trading

Models for trade trajectory optimization

- dependence of cost on scheduling decisions

Two conundrums of market impact

1. Buyer vs seller

   who pays impact to whom?

2. Impact vs alpha

   trade decision is not exogenous

   depends on previous price changes

   and on anticipated price changes

Conundrum #1: Buyer vs seller

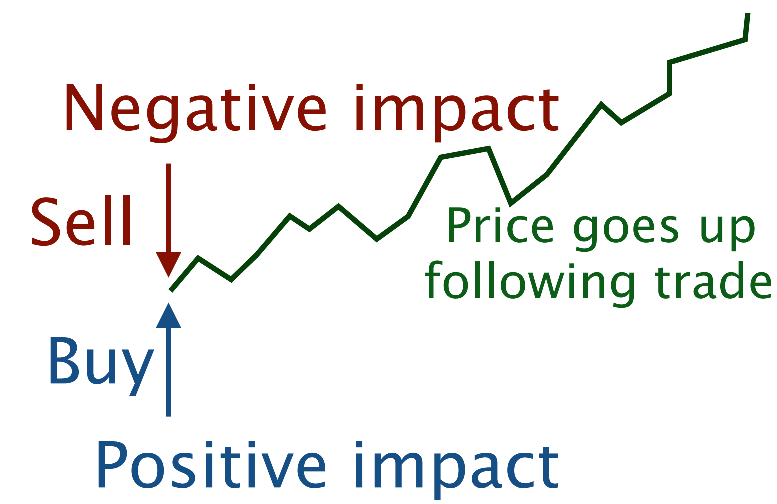

Every trade has two sides

- Which one pays market impact?

Answers

- "People like me" pay to "the market"
- More aggressive pays to less aggressive

Data sources

Public market data

- impact of aggressive (market) orders
- problem: algo executions can be 50% passive

Broker or internal data set

- client orders paying impact to market
- problem: in closed system, sum to zero

CME LDB data set (to 2012)

- trade volume tagged by "CTI code" (local/external)
- can demonstrate transfer to locals

Newedge

Research Snapshot

NEWEDGE PRIME BROKERAGE | INVESTOR RESEARCH

March 30, 2011

Robert Almgren
robertalmgren@quantitativebrokers.com

Galen Burghardt
galen.burghardt@newedge.com

A window into the world of futures market liquidity

The purpose of this snapshot is to call attention to an interesting data set maintained by the Chicago Mercantile Exchange (CME) that affords a unique insight into futures trading costs. As brokers, we use this data to help understand transactions costs and to keep them as low as possible for our clients.

The CME microstructure data allows us to conclude two things. First, those traders whom we traditionally think of as liquidity takers do in fact pay for access to the pool of liquidity afforded by the exchange. Second, the net price paid for liquidity is remarkably small given the size of the bid/ask spread. In this example, which highlights trading in 10-year Treasury note futures, we find that the average price paid by "liquidity takers" is about $3 per contract per round turn, while the value of the bid/ask spread is just over $15.

CME LDB data (no longer available)

Averages calculated over full trading days

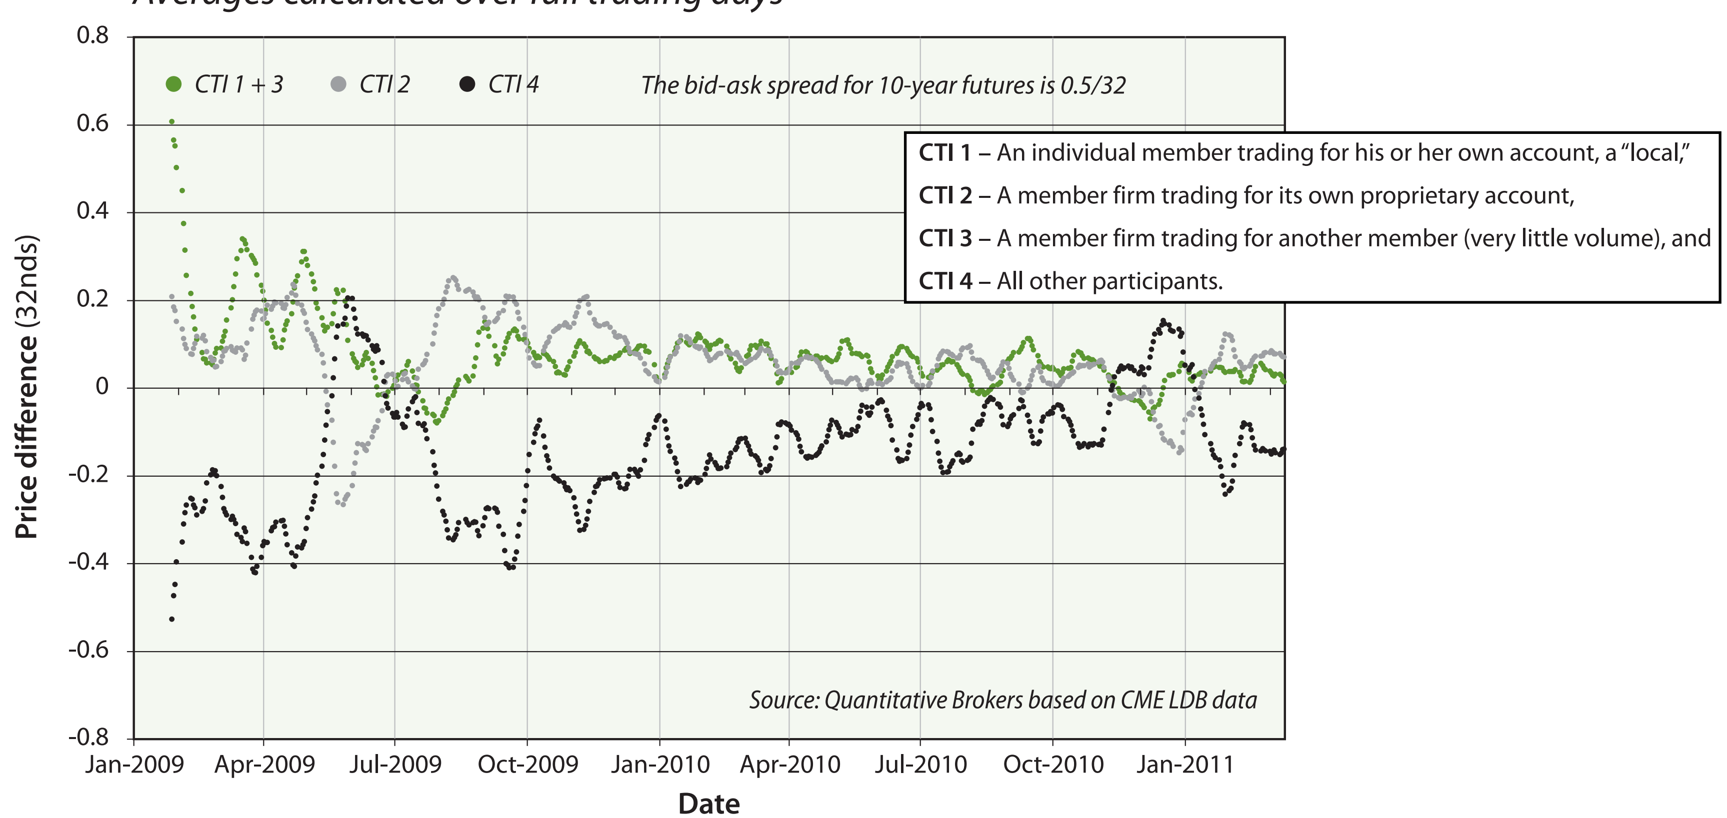

Conundrum #2: impact vs alpha

Decision to trade is never exogenous

Trader buys because expects price rise

- subsequent rise is impact or alpha?

Ideal study: send random orders

Must calibrate impact model for each trade style

- short-term alpha vs long term

Example: cross-impact

- correlation due to cross impact or correlated trading?

Example: serial correlation of trade sign

- buys followed by buys, sells by sells

Micro impact

Price change following market order execution
Only study you can do with public data

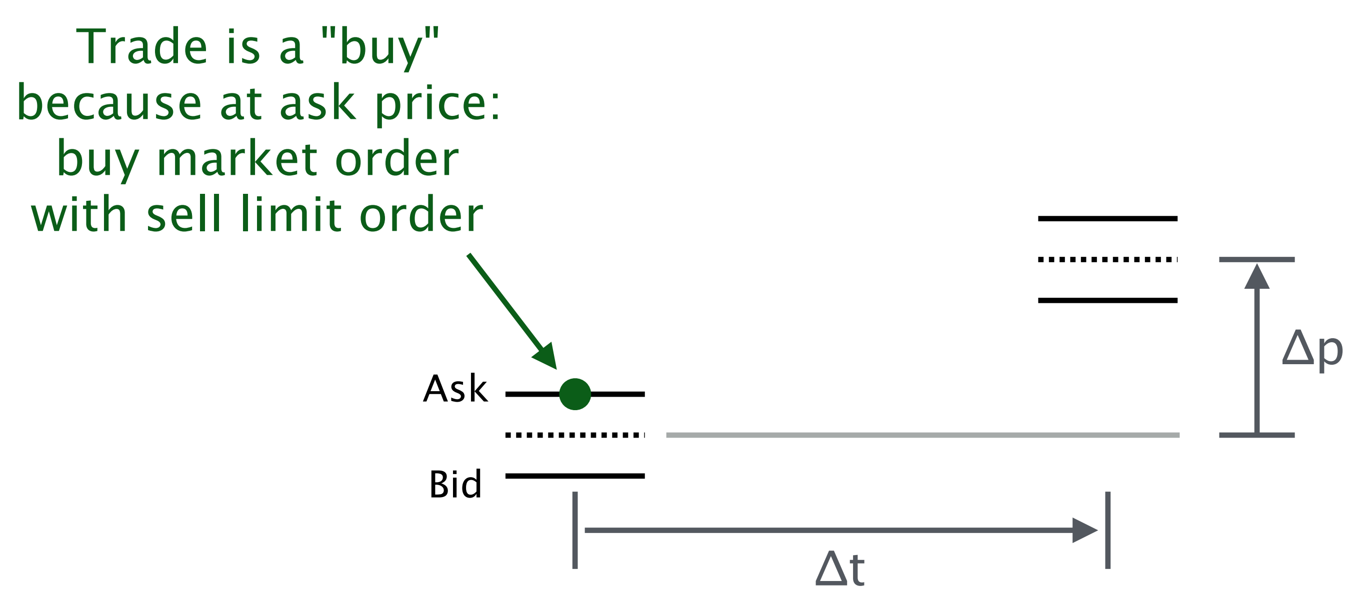

Market impact model:
Δp as function of
Δt and trade size

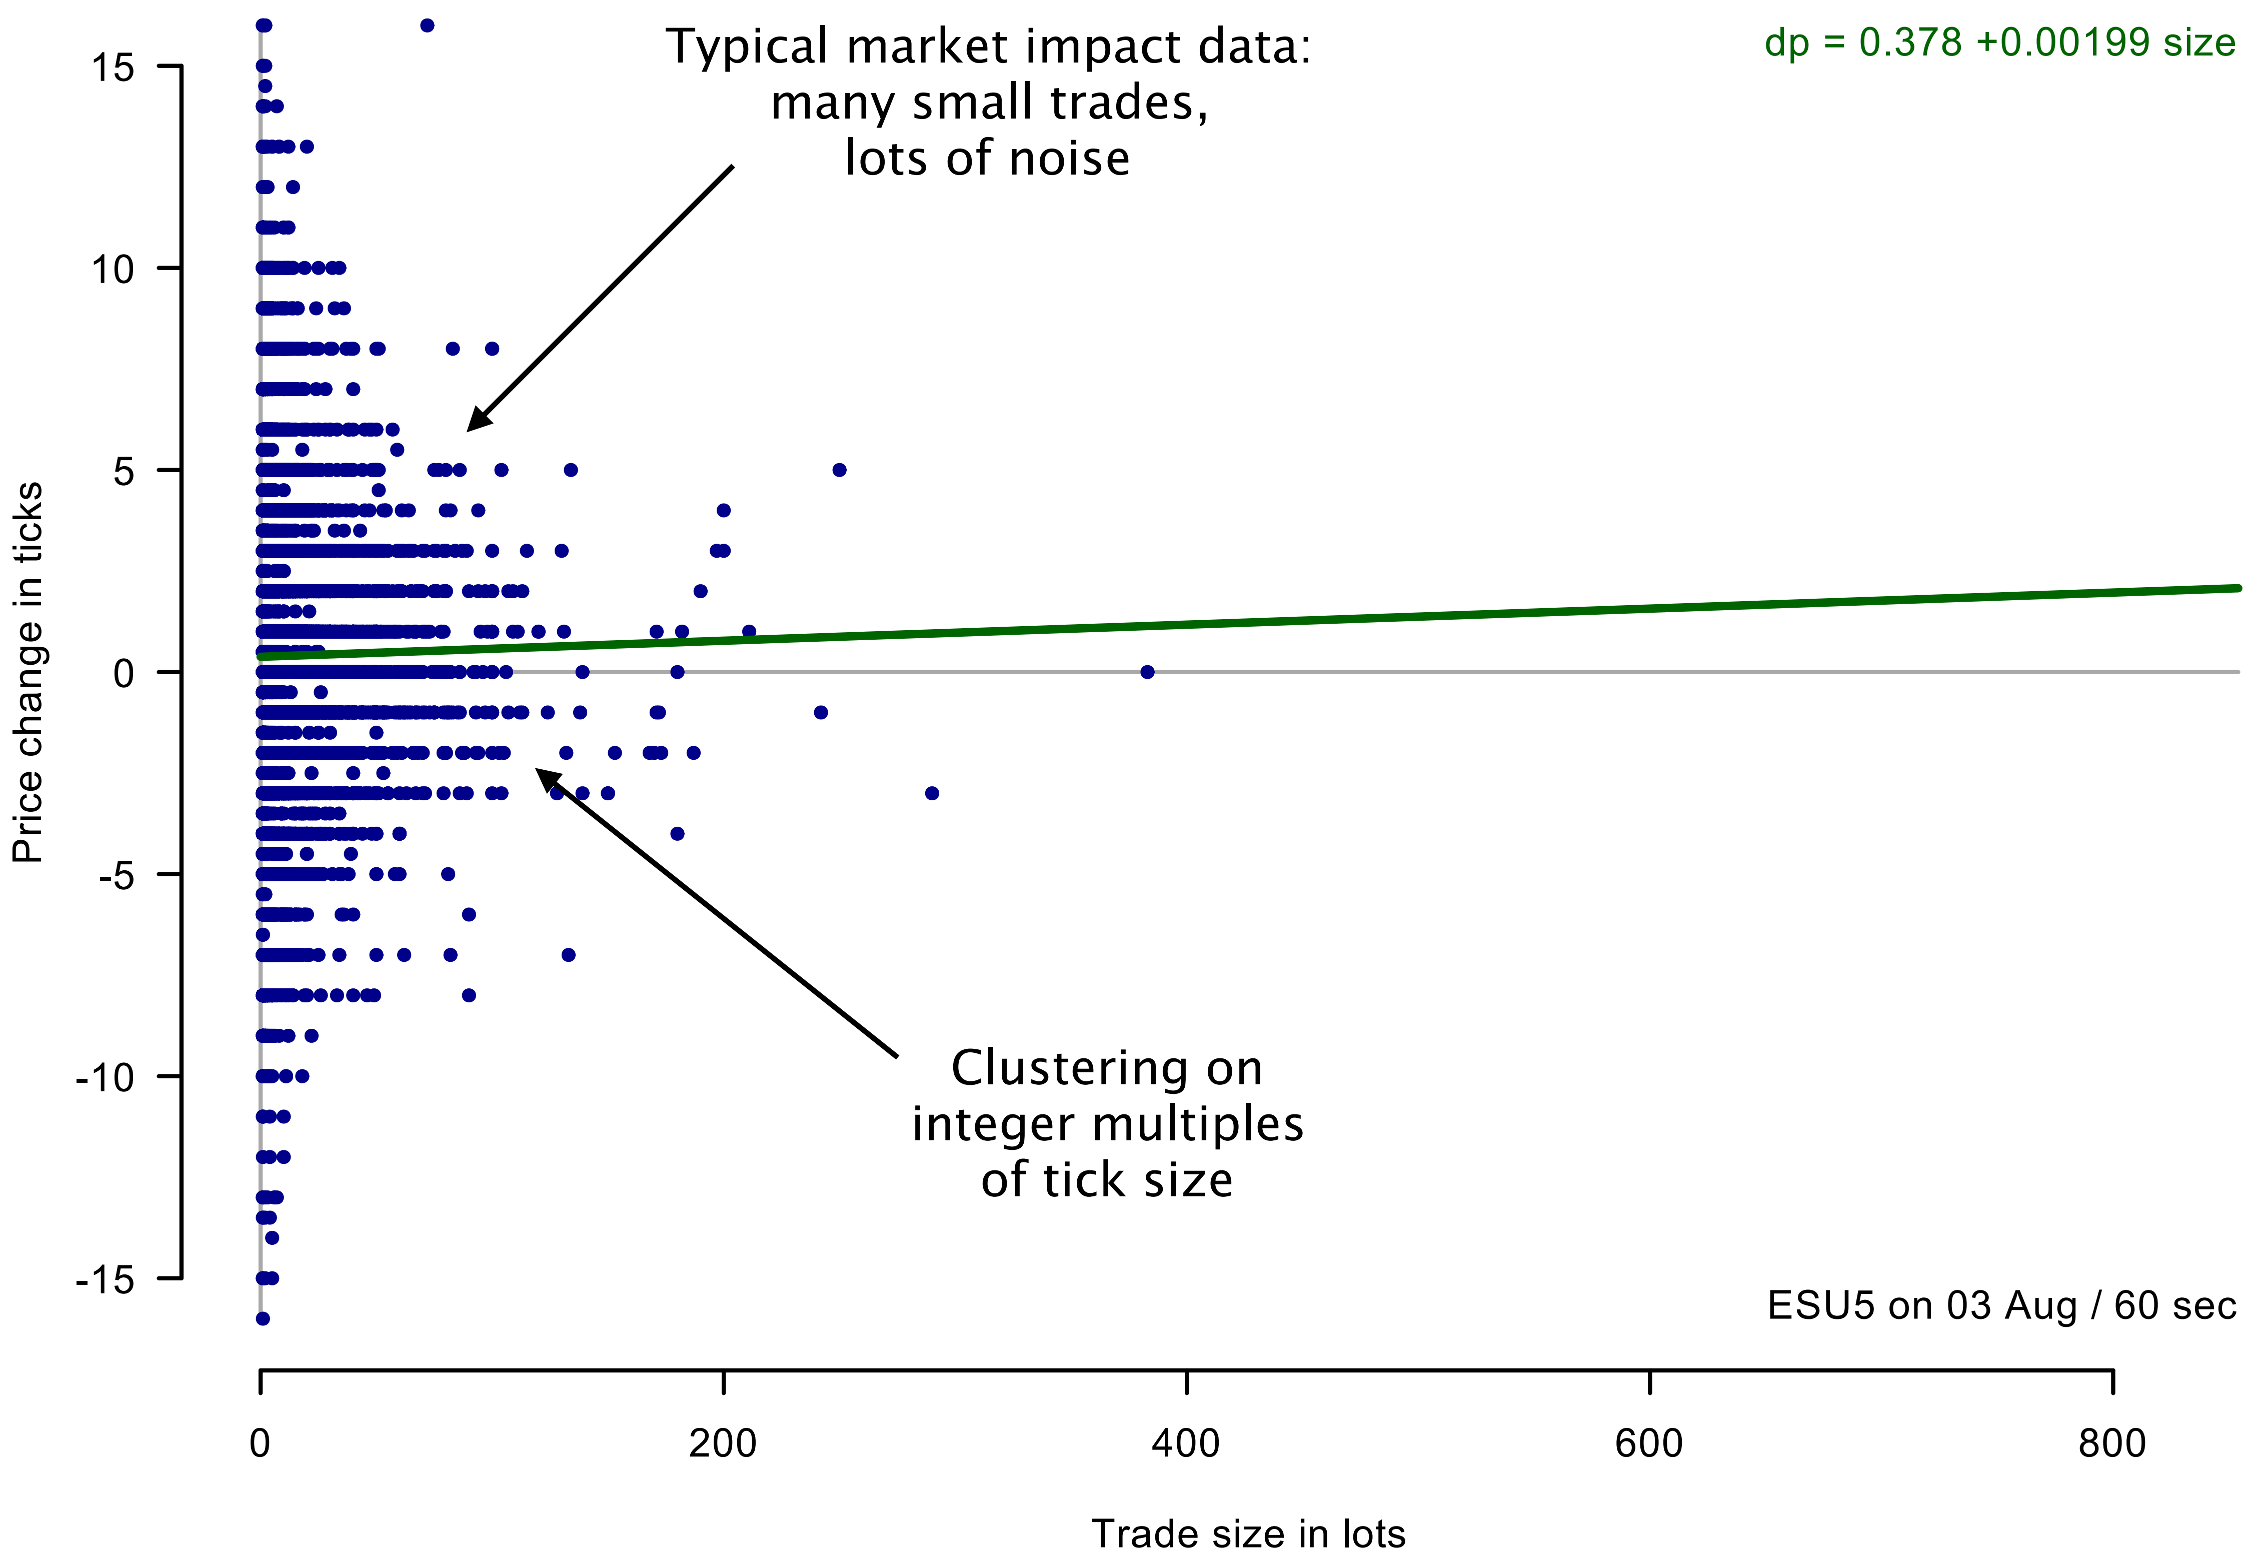
Typical market impact data:
many small trades,
lots of noise
dp = 0.378 +0.00199 size
Clustering on
integer multiples
of tick size
ESU5 on 03 Aug / 60 sec
Price change in ticks
Trade size in lots
15
10
5
0
-5
-10
-15
0
200
400
600
800

Same with log scale for trade size

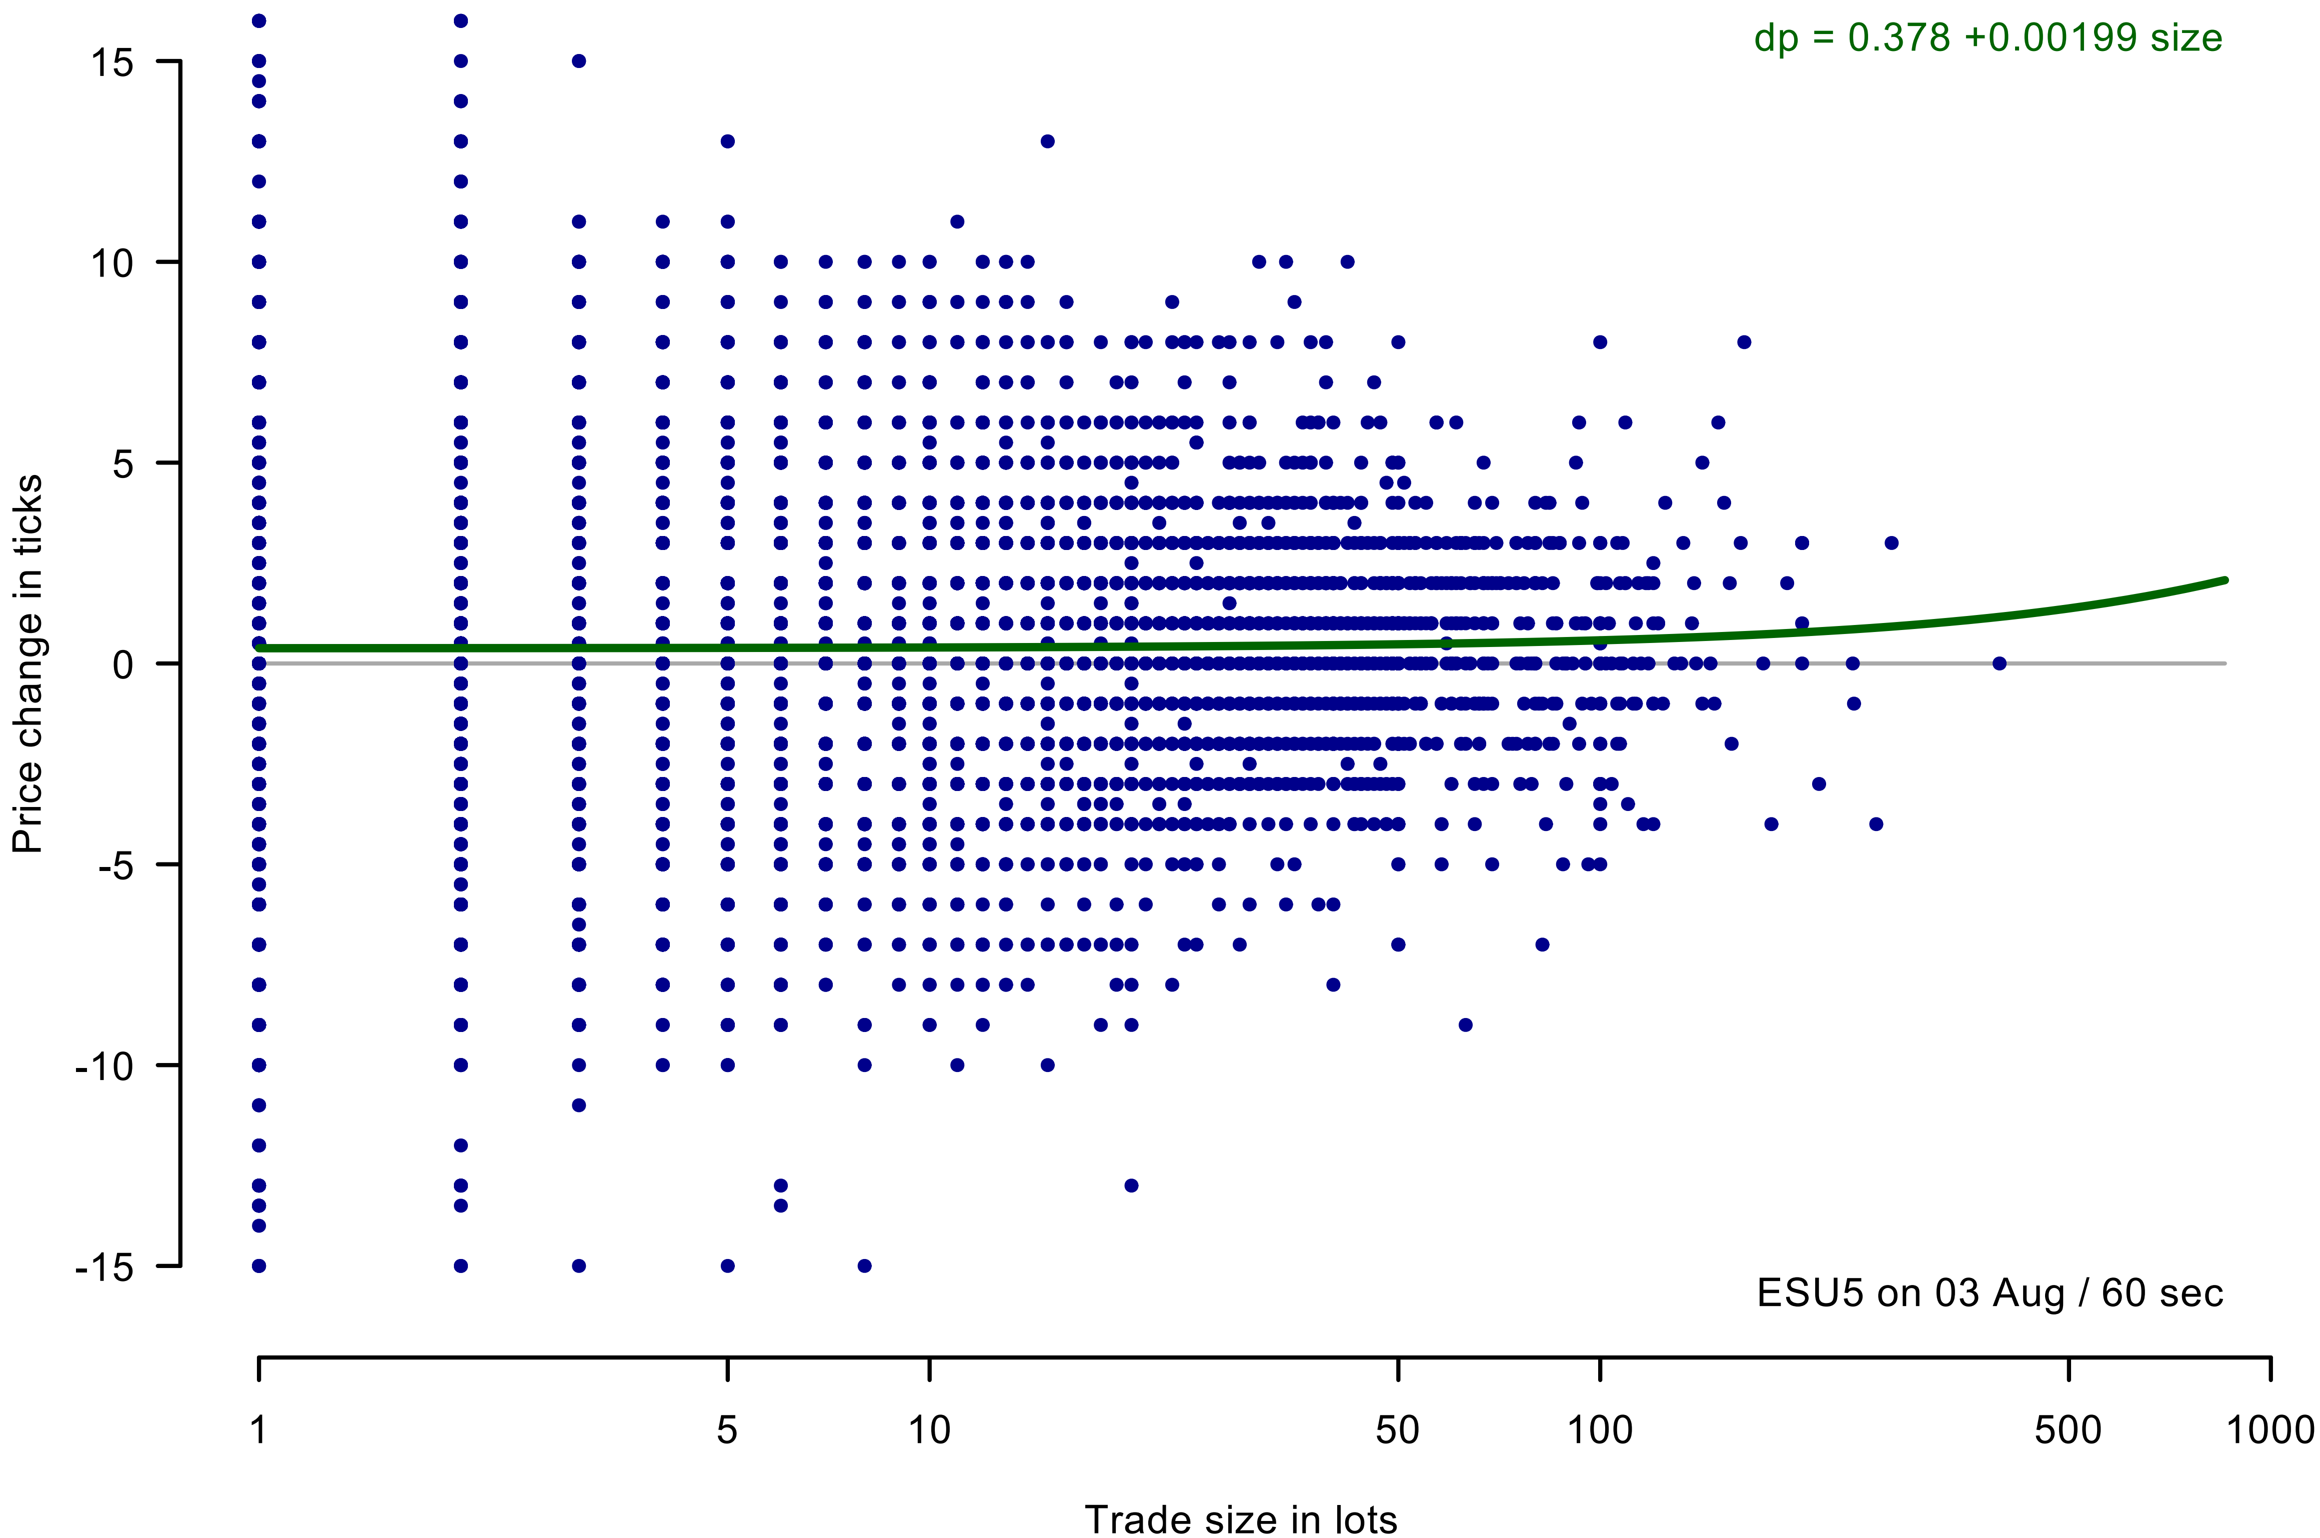

Challenges in micro impact

Buy/sell classification is arbitrary

- legitimate study, but may not be what you want

Market order may depend on quote sizes

- microprice (quote imbalance) is common signal

Market orders are serially correlated

- impact may be due to earlier and later orders
- cause may be slicing of larger orders, or
- trade decisions reacting to trade activity

$$p_{\mathrm{micro}} = \frac{q_{\mathrm{bid}}\, p_{\mathrm{ask}} + q_{\mathrm{ask}}\, p_{\mathrm{bid}}}{q_{\mathrm{bid}} + q_{\mathrm{ask}}}$$

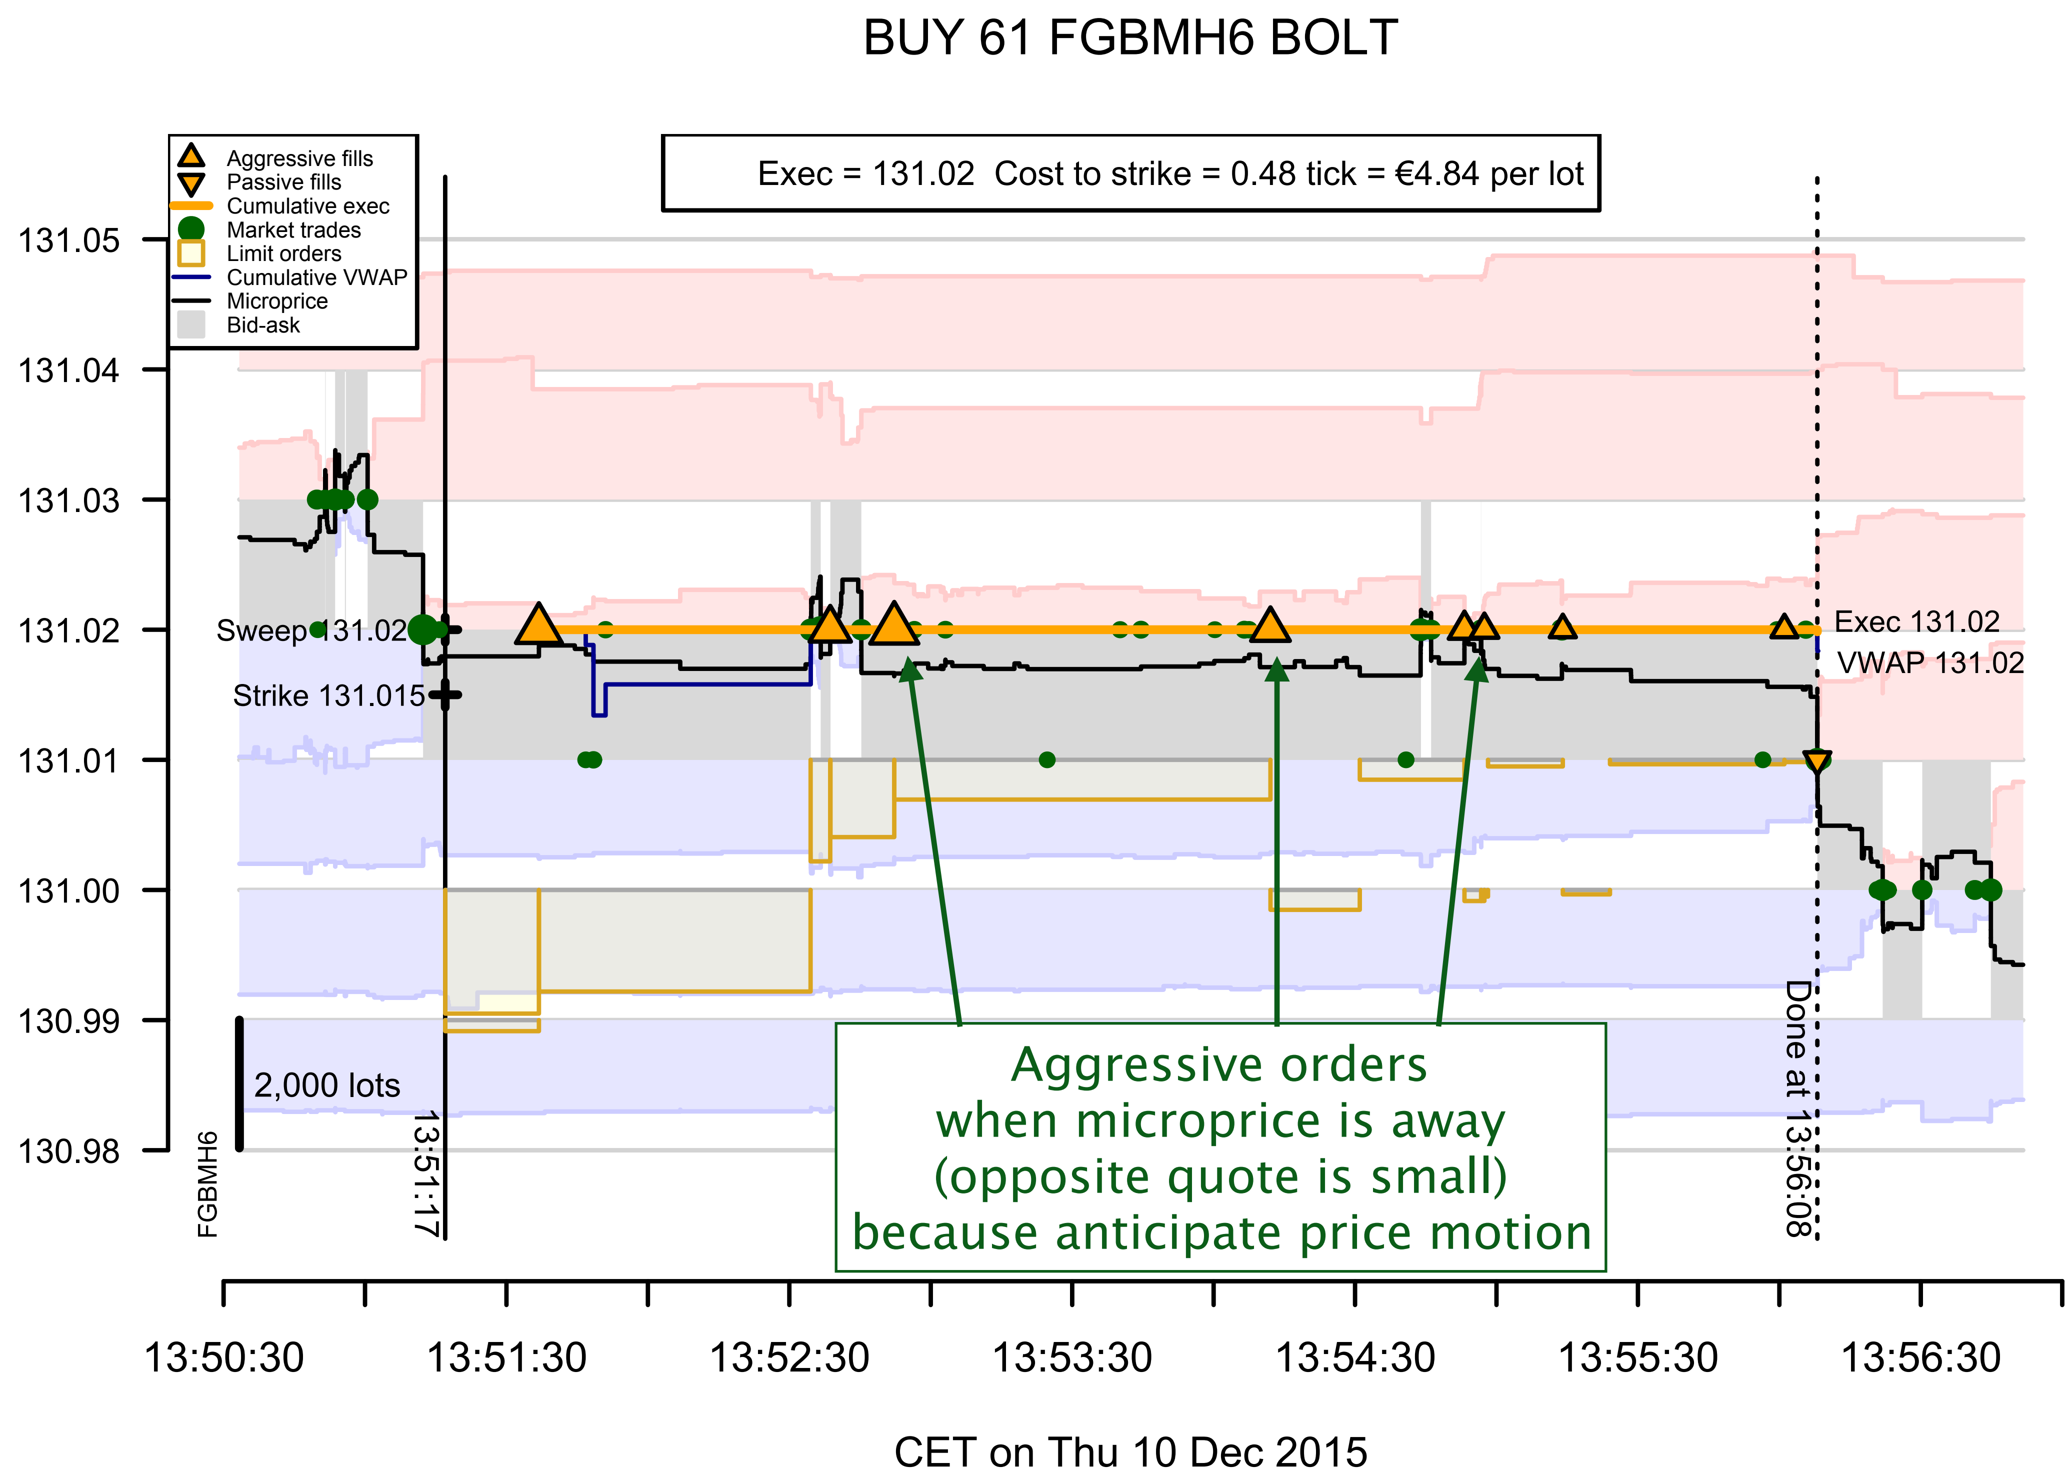

SIAM J. FINANCIAL MATH.
Vol. 4, pp. 1–25

Price Dynamics in a Markovian Limit Order Market*

Rama Cont† and Adrien de Larrard†

Abstract. We propose a simple stochastic model for the dynamics of a limit order book, in which arrivals of market orders, limit orders, and order cancellations are described in terms of a Markovian queueing system. Price dynamics are endogenous and result from the execution of market orders against outstanding limit orders. Through its analytical tractability, the model allows us to obtain analytical expressions for various quantities of interest, such as the distribution of the duration between price changes, the distribution and autocorrelation of price changes, and the probability of an upward move in the price, *conditional* on the state of the order book. We study the diffusion limit of the price process and express the volatility of price changes in terms of parameters describing the arrival rates of buy and sell orders and cancellations. These analytical results provide some insight into the relation between order flow and price dynamics in limit order markets.

Key words. limit order book, market microstructure, queueing, diffusion limit, high-frequency data, liquidity, duration analysis, point process

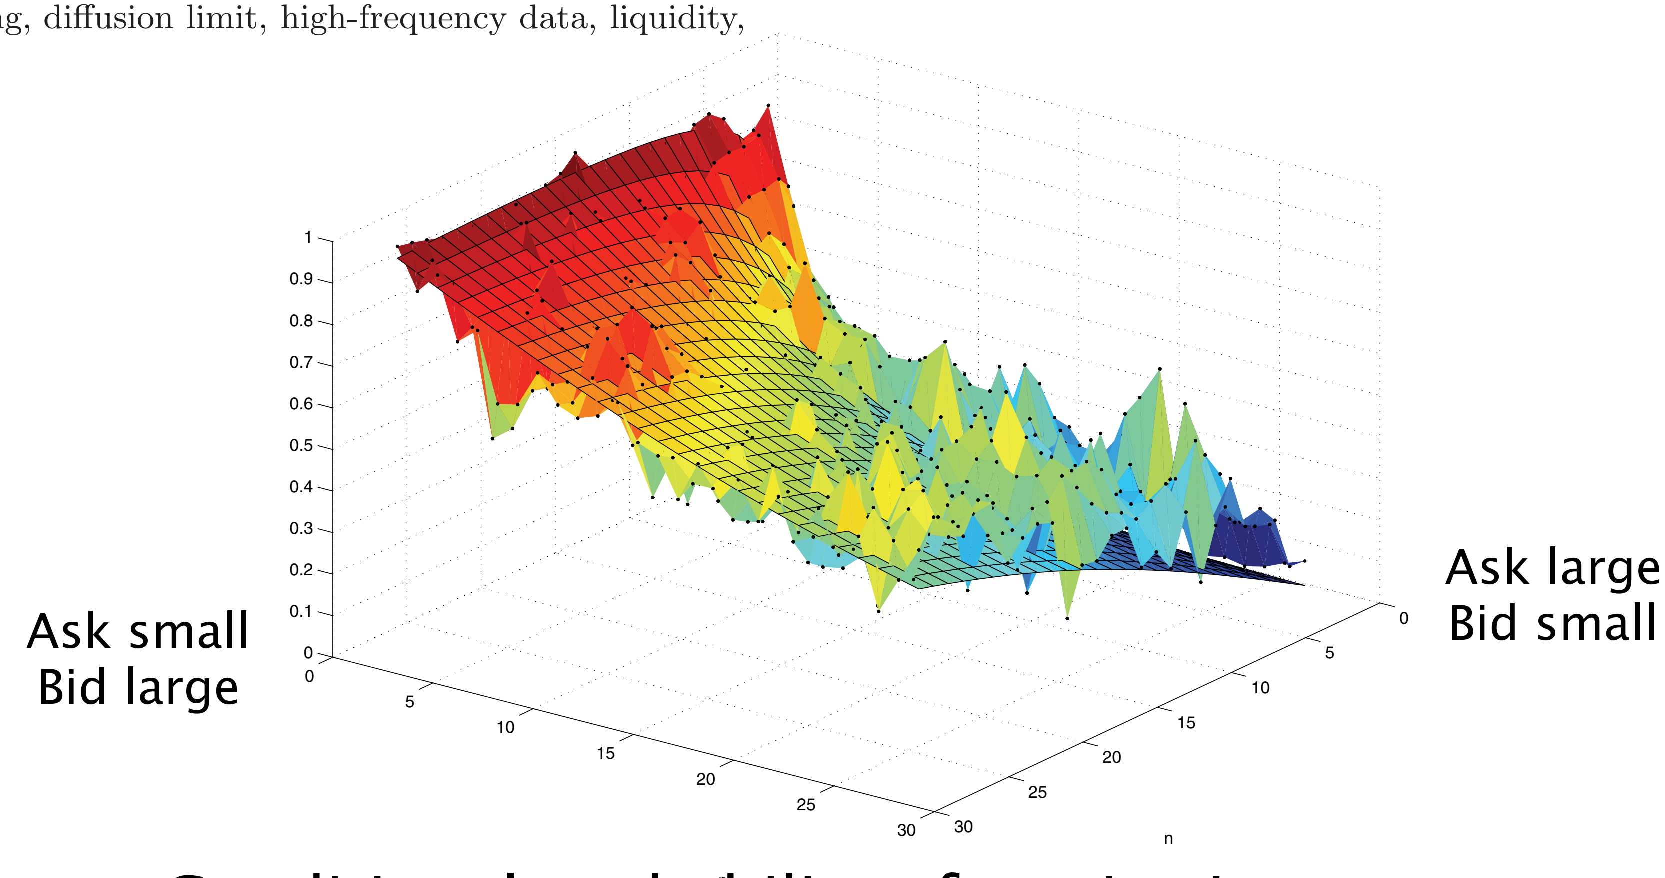

Conditional probability of a price increase

Figure 5. *Conditional probability of a price increase, as a function of the bid and ask queue sizes, compared with empirical transition frequencies for Citigroup stock price tick-by-tick data on June 26th, 2008.*

Modern models: "propagators"

RESEARCH PAPER
quant.iop.org

QUANTITATIVE FINANCE VOLUME 4 (2004) 176–190
INSTITUTE OF PHYSICS PUBLISHING

Fluctuations and response in financial markets: the subtle nature of 'random' price changes

Jean-Philippe Bouchaud[1,2], Yuval Gefen[3], Marc Potters[2] and Matthieu Wyart[1]

[1] Commissariat à l'Energie Atomique, Orme des Merisiers, 91191 Gif-sur-Yvette Cedex, France
[2] Science and Finance, Capital Fund Management, 109–111 rue Victor Hugo, 92 532 Levallois Cedex, France
[3] Condensed Matter Physics Department, Weizmann Institute of Science, 76 100 Rehovot, Israel

More precisely, one can consider the following correlation function:

$$\mathcal{C}_0(\ell) = \langle \varepsilon_{n+\ell}\varepsilon_n \rangle - \langle \varepsilon_n \rangle^2. \qquad (6)$$

If trades were random, one should observe that $\mathcal{C}_0(\ell)$ decays to zero beyond a few trades. Surprisingly, this is not what happens: on the contrary, $\mathcal{C}_0(\ell)$ is strong and decays very slowly toward zero, as an inverse power law of ℓ (see figure 9):

$$\mathcal{C}_0(\ell) \simeq \frac{C_0}{\ell^\gamma}, \qquad (\ell \geqslant 1). \qquad (7)$$

In order to better understand the impact of trading on price changes, one can study the following *response function* $\mathcal{R}(\ell)$, defined as

$$\mathcal{R}(\ell) = \langle (p_{n+\ell} - p_n)\varepsilon_n \rangle, \qquad (3)$$

where ε_n is the sign of the nth trade, introduced in section 2.1. The quantity $\mathcal{R}(\ell)$ measures how much, on average, the price moves up conditioned to a buy order at time zero (or a sell order moves the price down) a time ℓ later.

Price motion has no serial correlation, even though is response to correlated order flow. Other traders *anticipate* future orders.

Macro impact

Need to know "parent order"
Plot slippage vs size
Fit linear or nonlinear model

Cost model

Inputs:

X = executed order size

B = benchmark price, bid-ask midpoint at start

P = average executed price

C = P - B = trade cost or slippage (for buy order)

  = -(P - B) for sell order

Model C as function of X: C = f(X)

This structure takes no account of
how the order is executed
or over what time horizon.
No use for optimizing execution!

Nondimensionalization

$$\frac{C}{\sigma} = f\left(\frac{X}{V}\right)$$

V = daily volume (actual, average, or forecast)
σ = daily volatility

Idea: Measure your order relative to
what the market is doing anyway
Lets you compare different assets and different days
(with widely varying volume and volatility) in same model

"Trading 1% daily volume costs 5% of daily volatility"

Structure of f(x)

$$f(x) = a + bx$$ Linear

$$f(x) = a + bx^k$$ Nonlinear

Minimize sum of squares error to order data

j indexes orders
a, b, k are universal

$$\frac{C_j}{\sigma_j} = a + b\left(\frac{X_j}{V_j}\right)^k + \epsilon_j, \quad \epsilon_j \text{ i.i.d.}$$

$$\min_{a,b,k} \sum_j \left(a + b\left(\frac{X_j}{V_j}\right)^k - \frac{C_j}{\sigma_j}\right)^2$$

SP500 (ES) 2015 (unscaled)

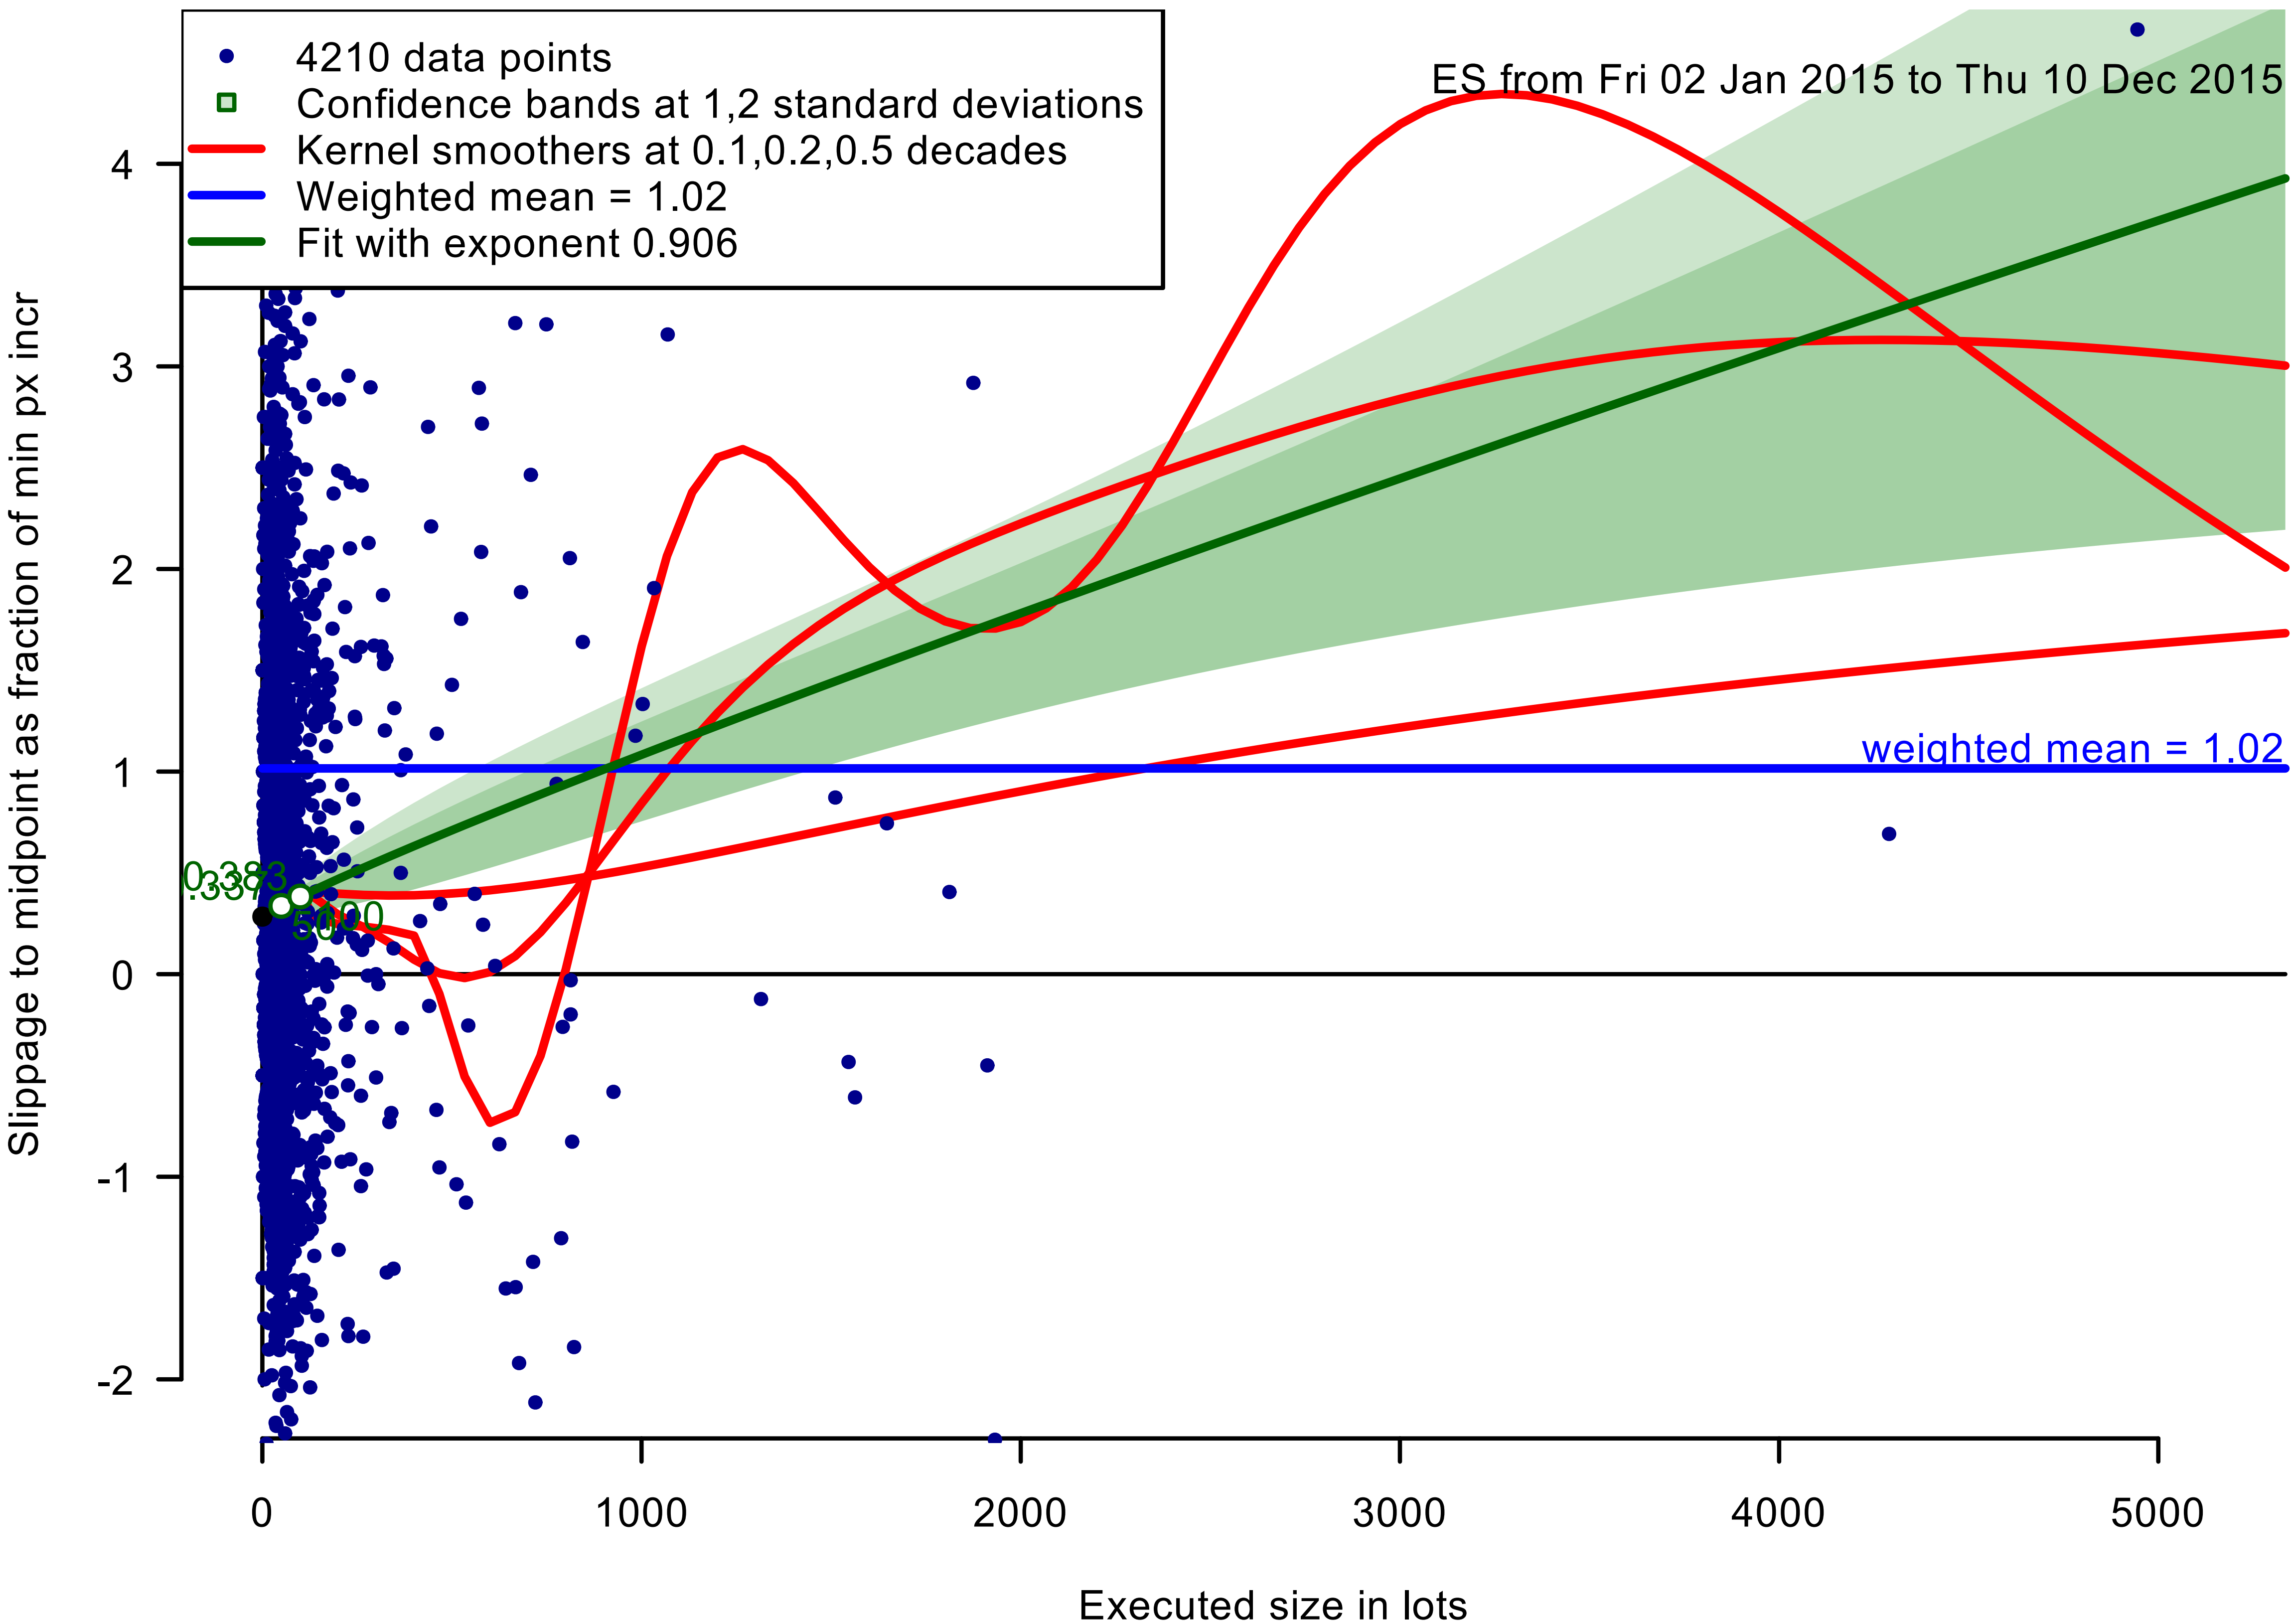

SP500 (ES) 2015 (unscaled)

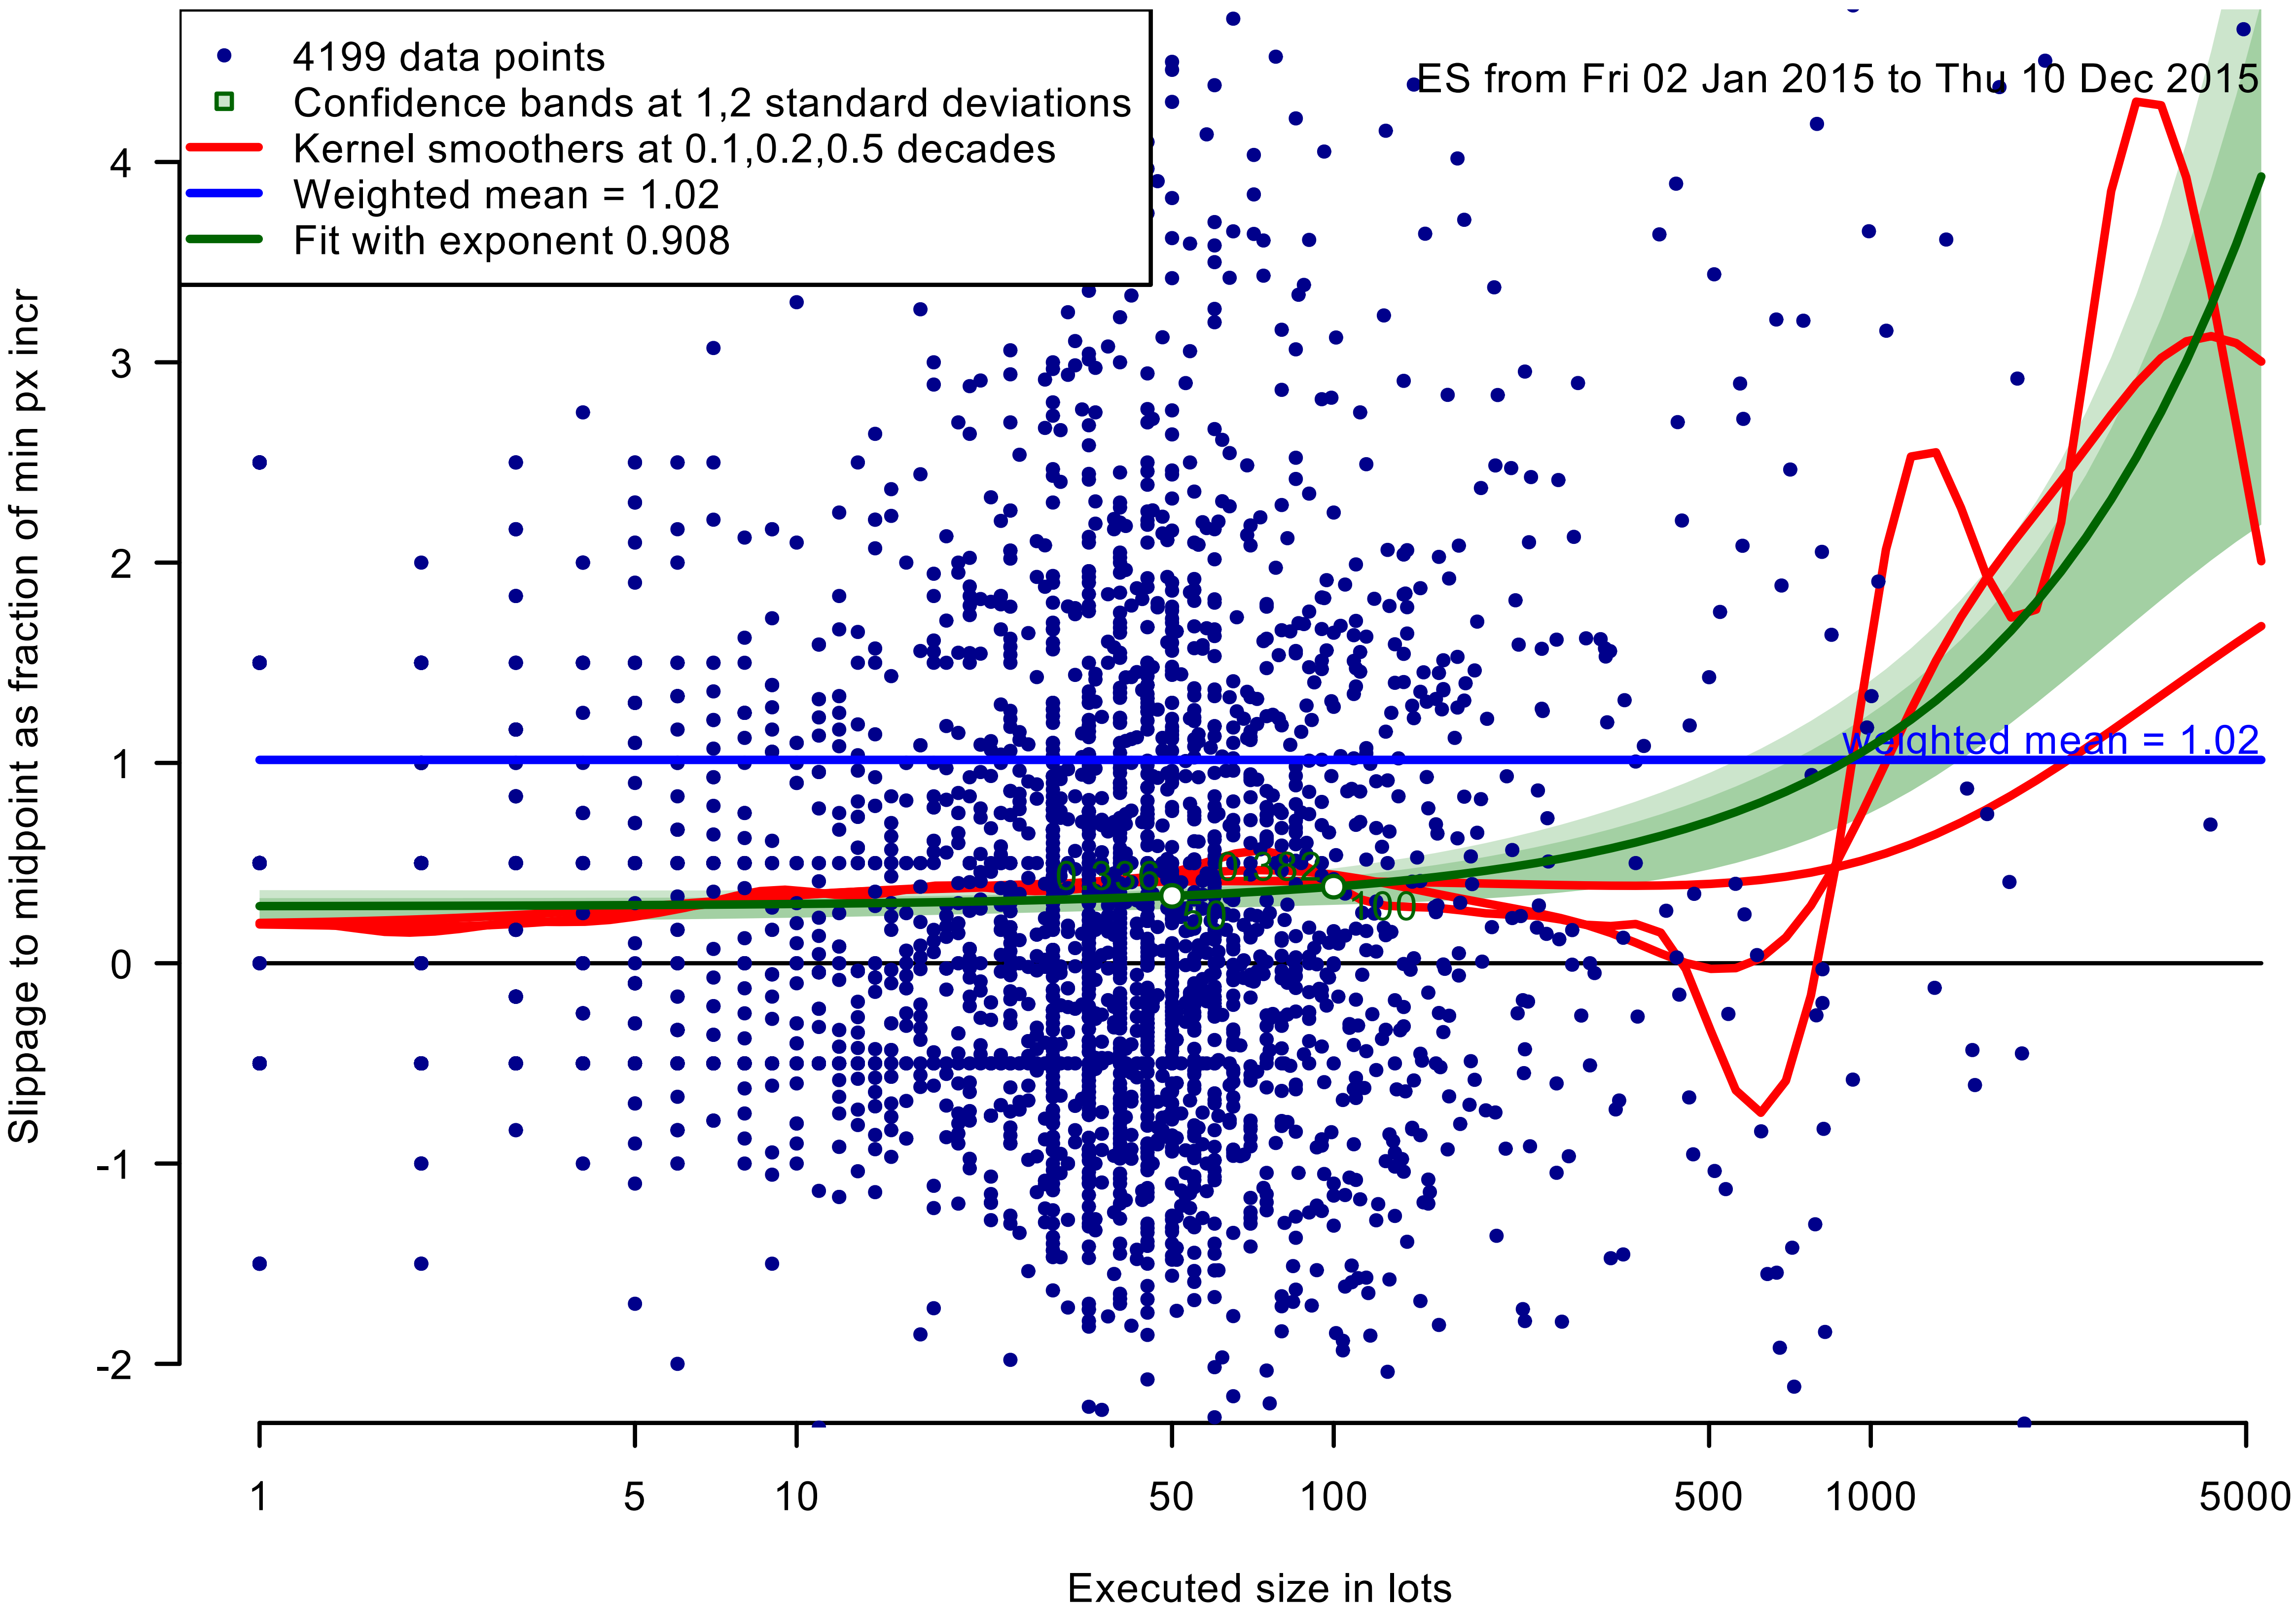

Crude Oil (CL) 2015 (unscaled)

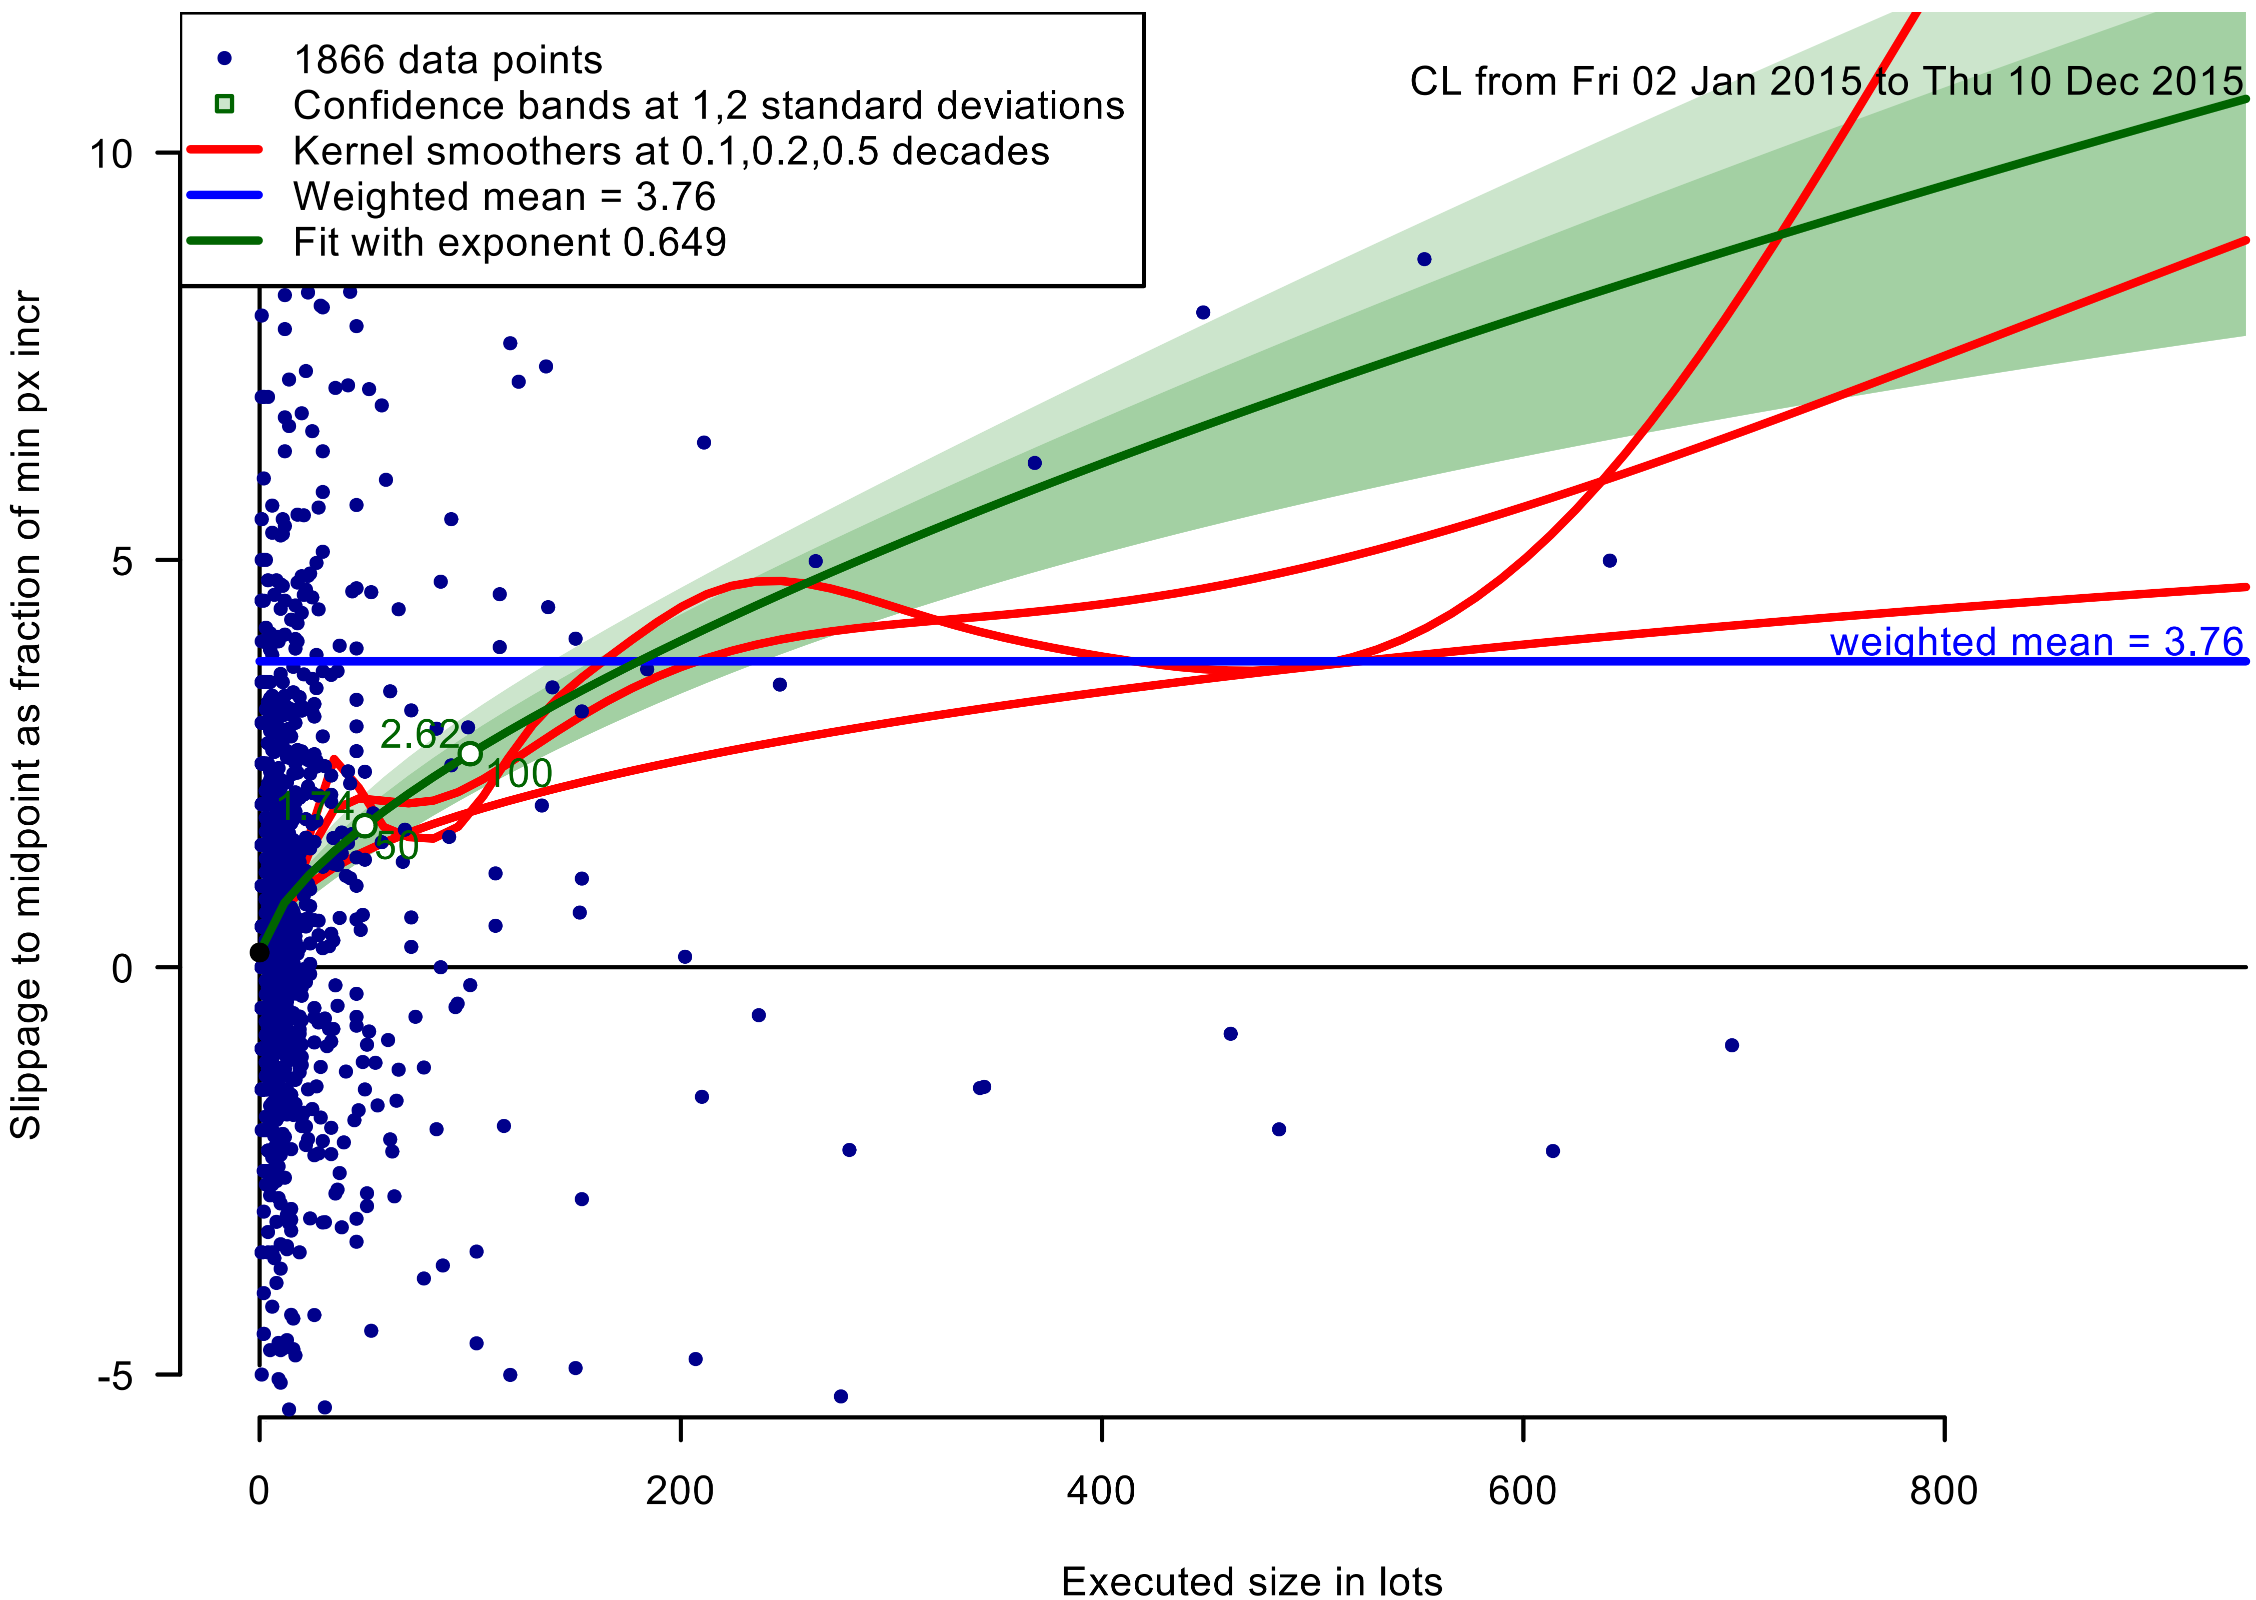

Crude Oil (CL) 2015 (unscaled)

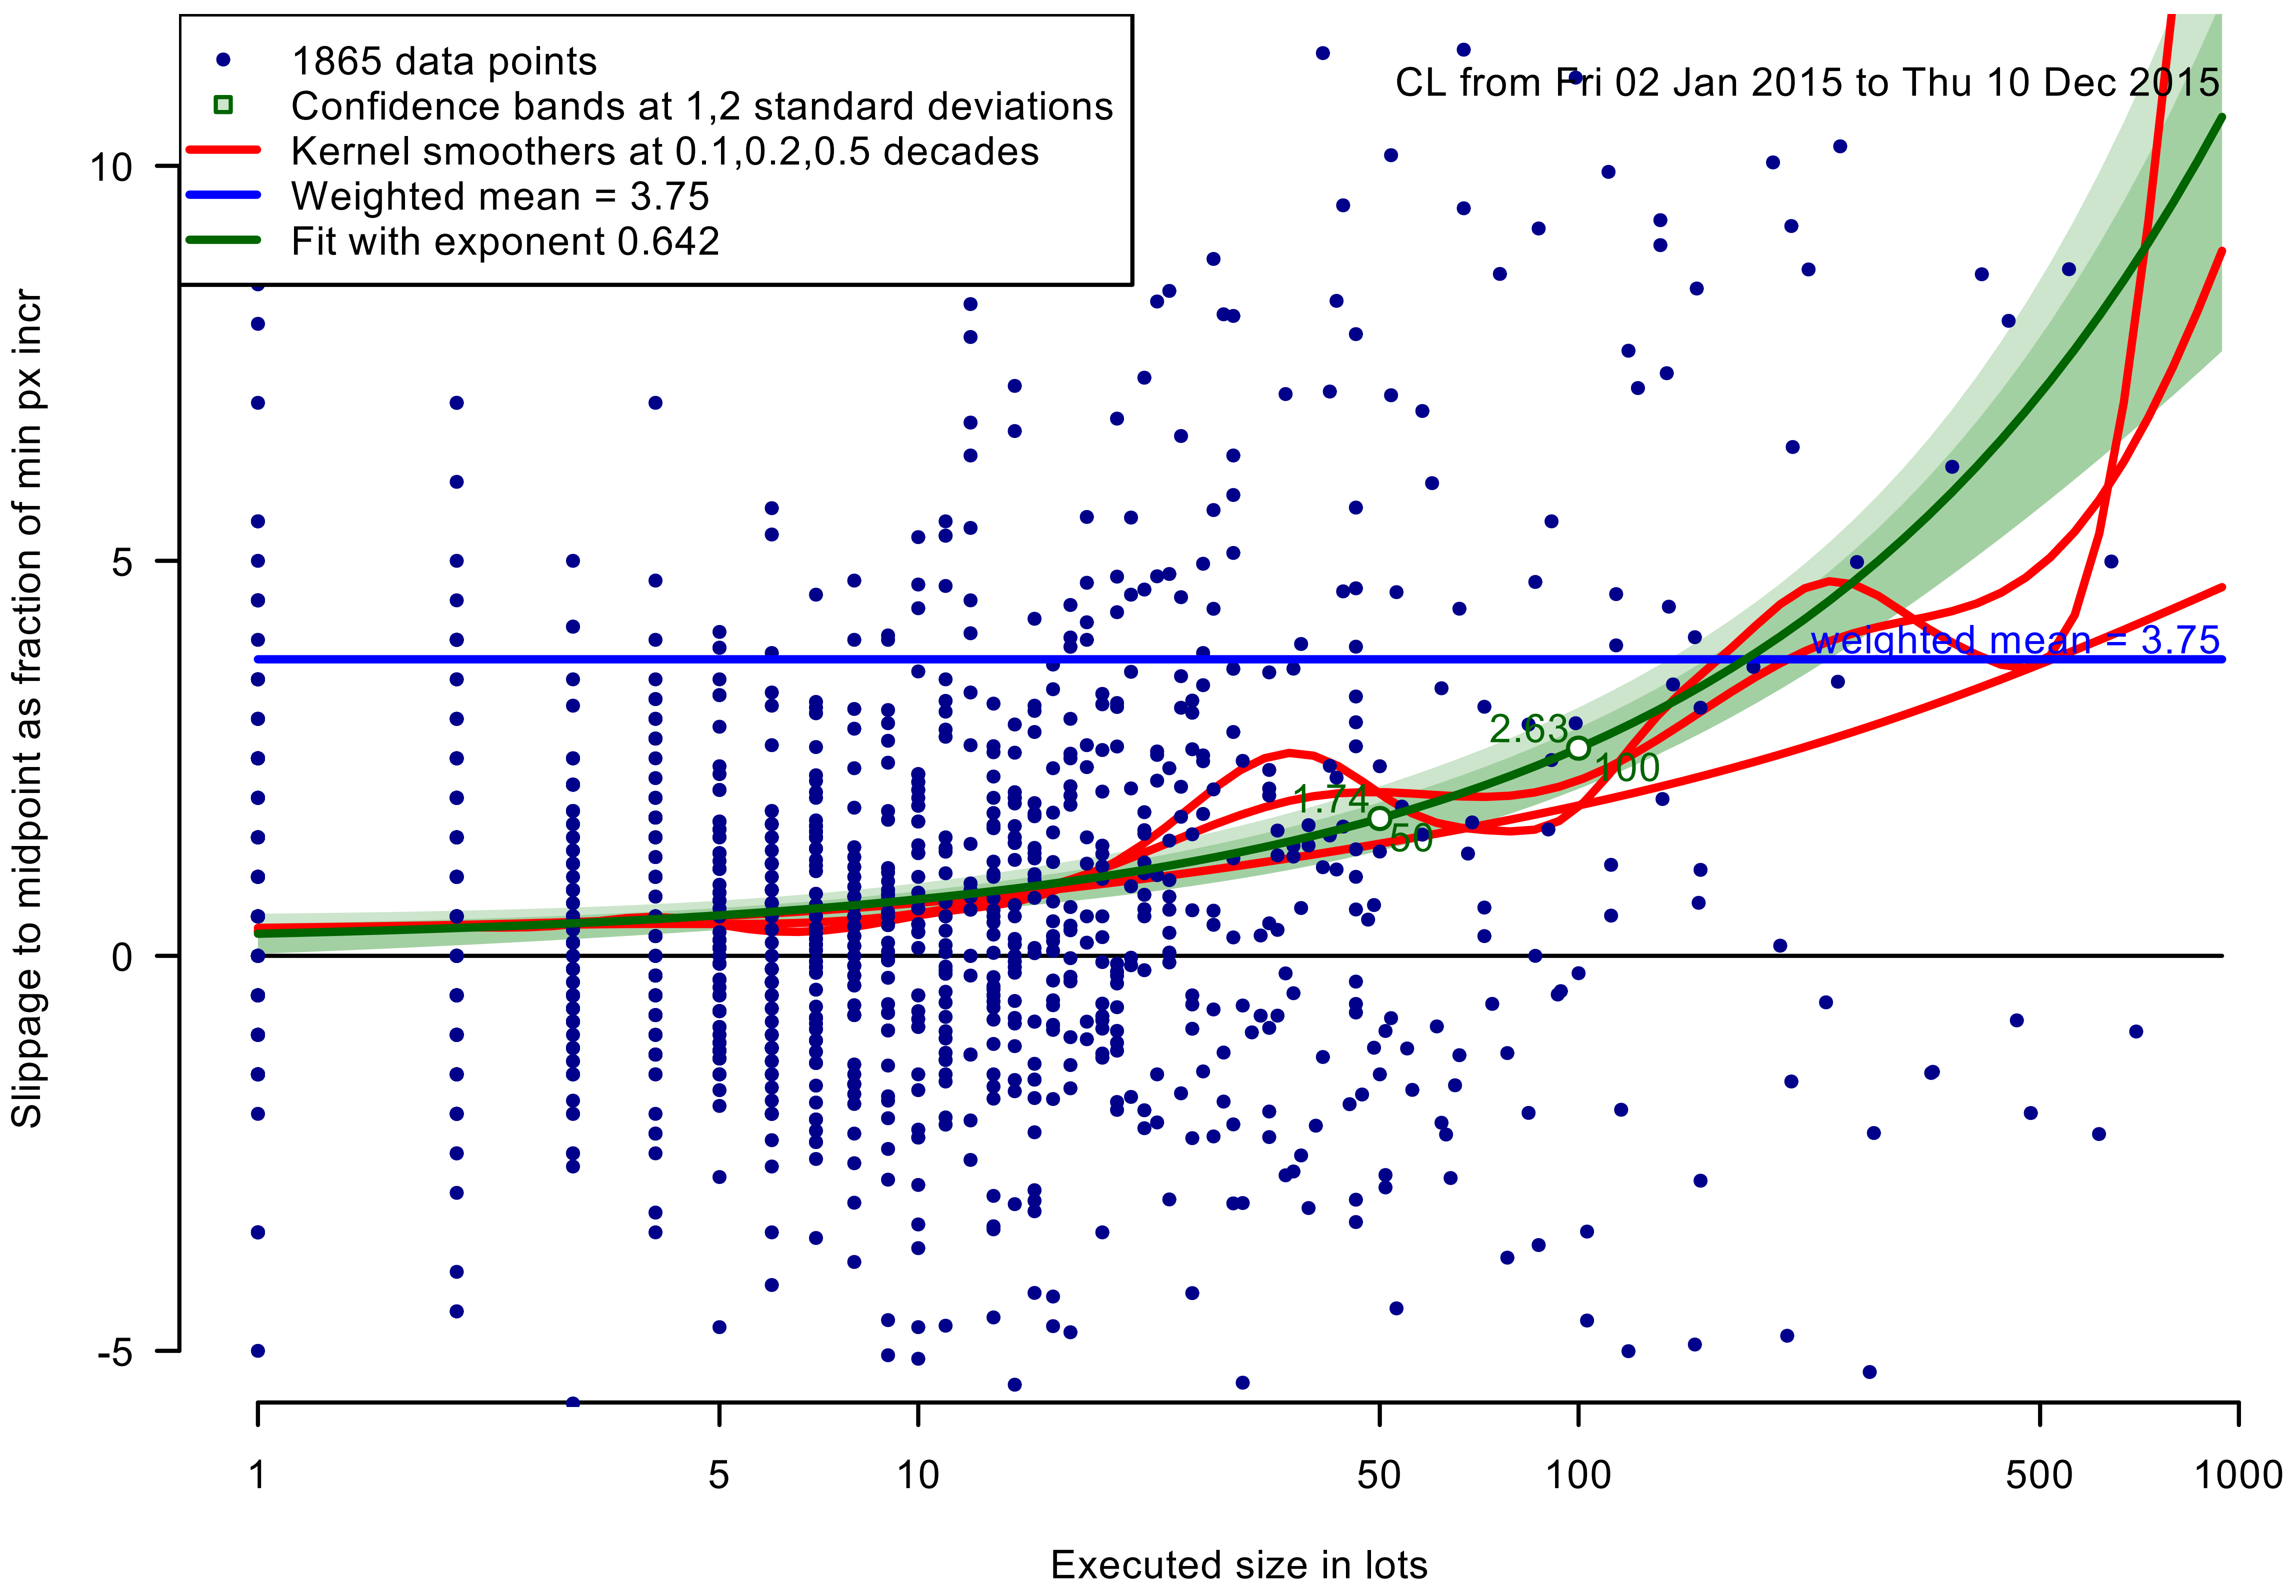

Scaled fit for several energy products

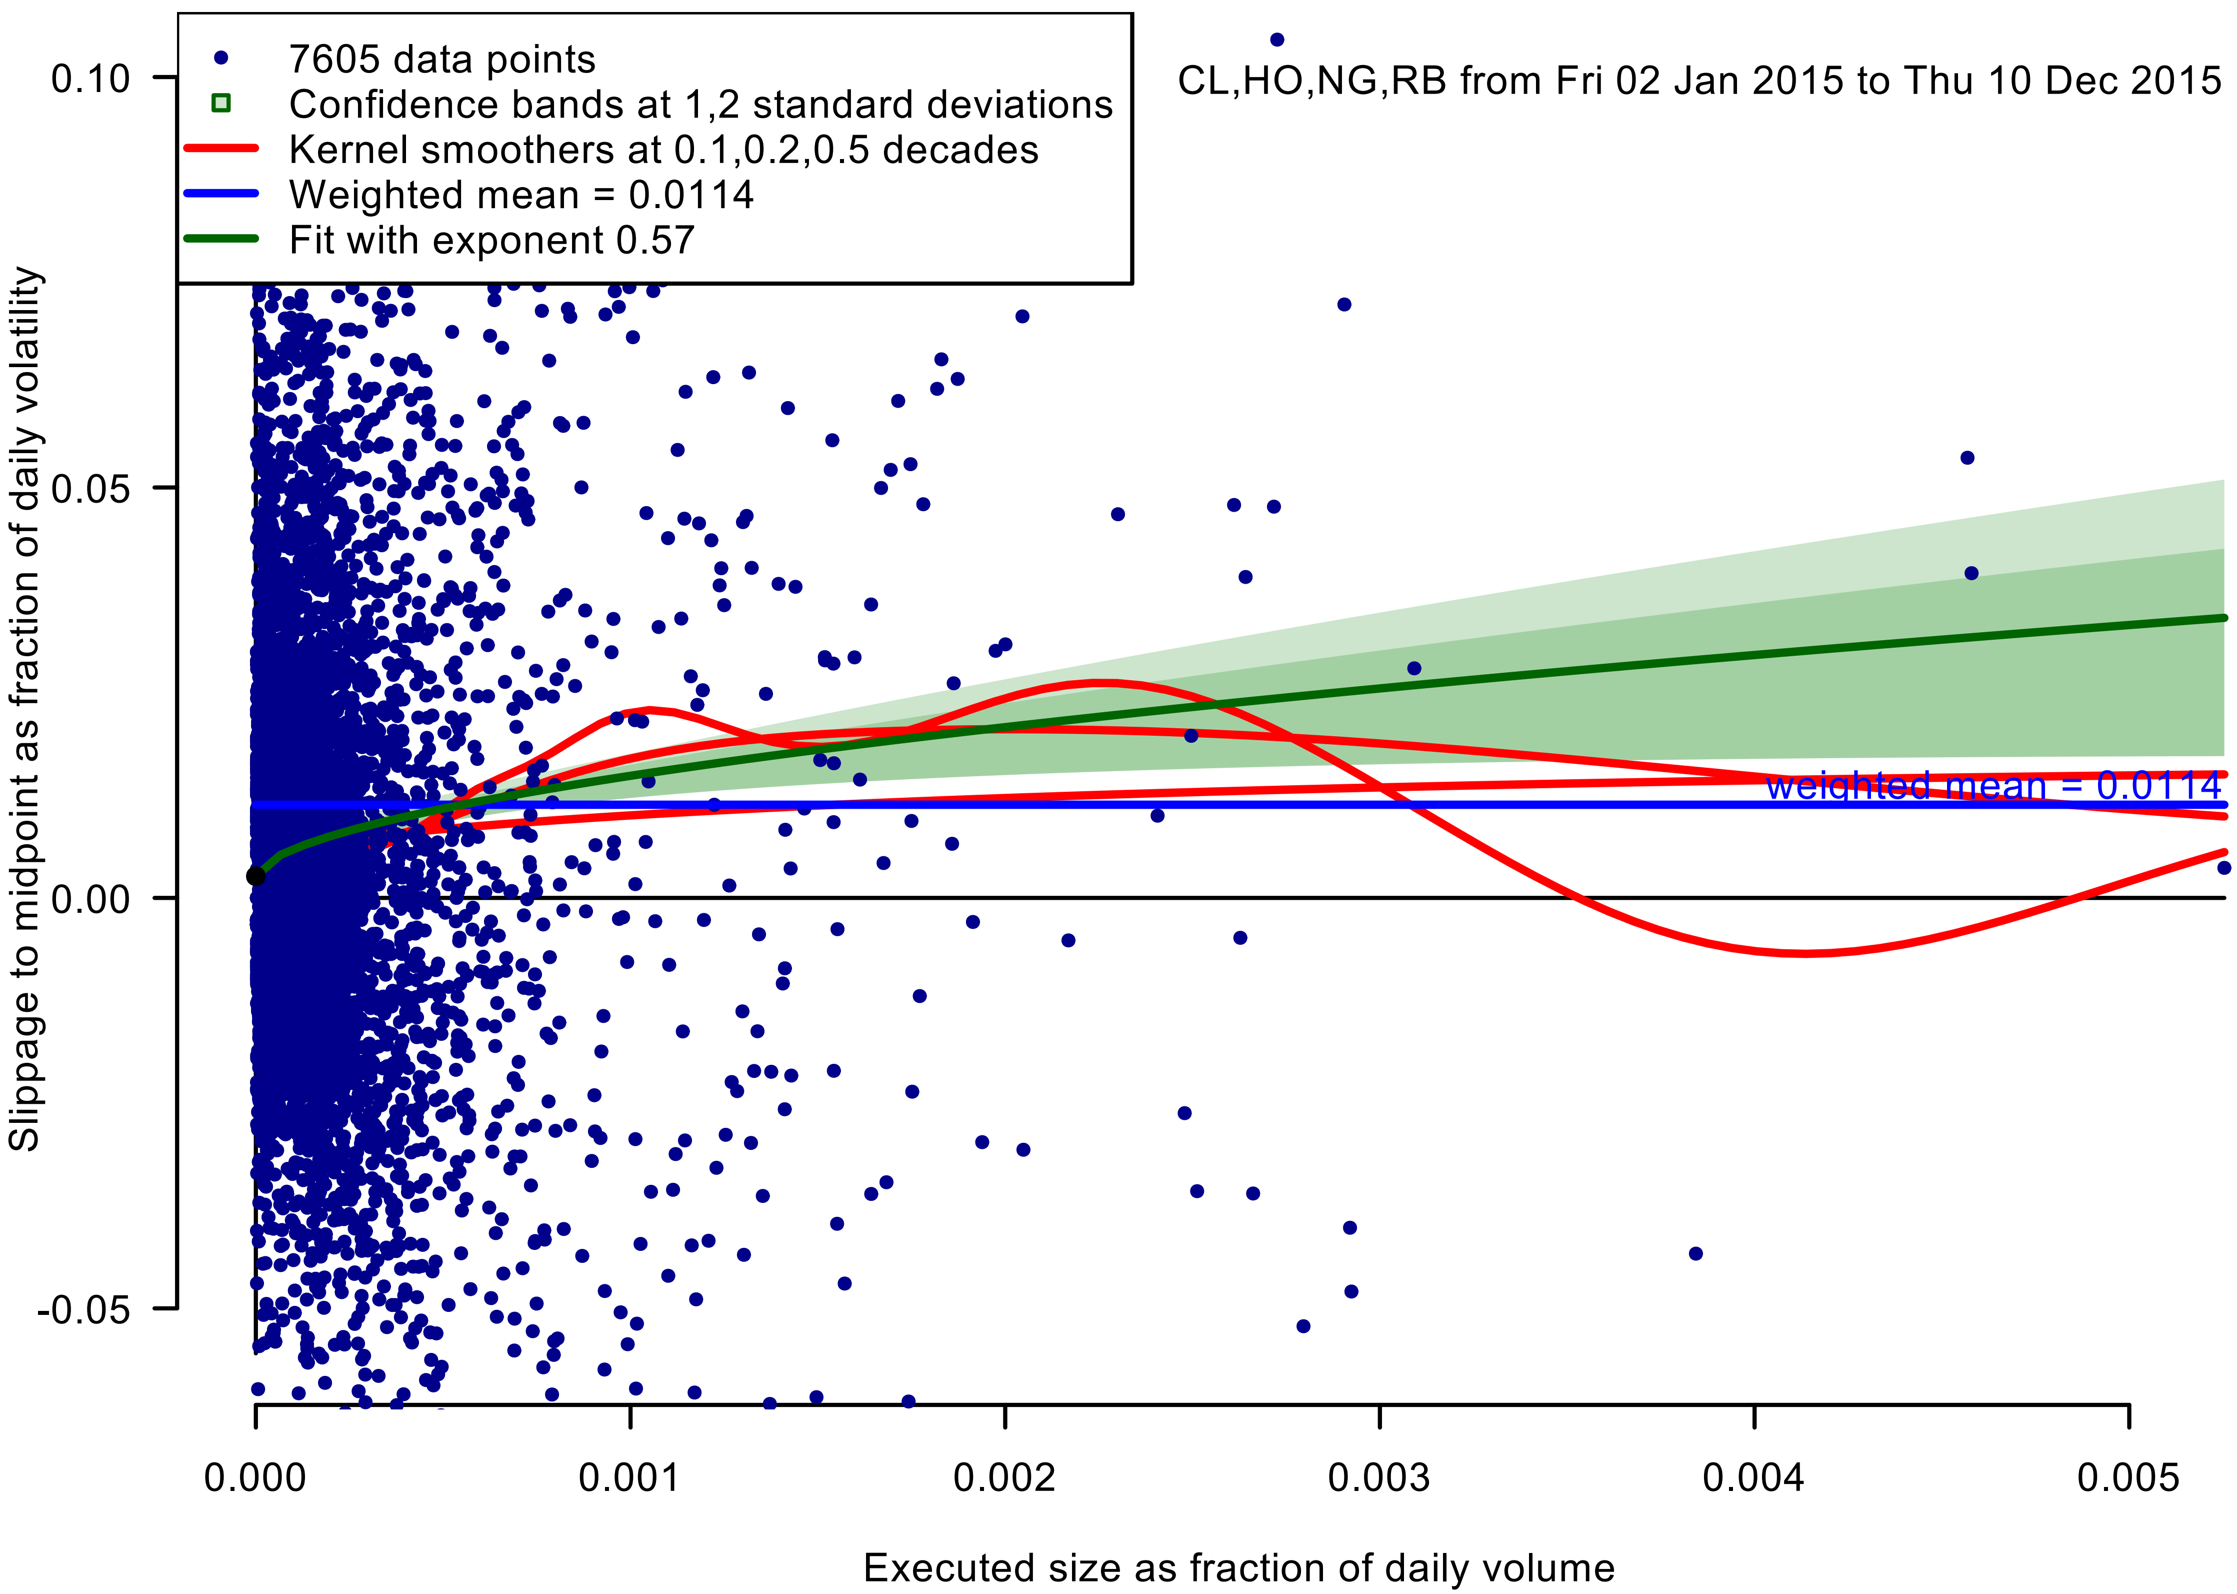

Advantages of this model (parent order level):

- simple to do
- simple to interpret
- gives immediate useful results for cost estimation

Disadvantages of this model

- not useful for order scheduling or optimization
- no microscopic description of mechanism

Caveats

- some orders may be cancelled based on market moves
  - solution: restrict sample to fully executed orders
- different strategies have different short term alpha
  - solution: results are client-specific and strategy-specific

Direct Estimation of Equity Market Impact*

Robert Almgren,[†] Chee Thum,[‡]
Emmanuel Hauptmann,[‡] and Hong Li[‡]

Risk, July 2005.

[†]University of Toronto Departments of Mathematics and Computer Science, and Citigroup Global Quantitative Research; Robert.Almgren@utoronto.ca.
[‡]Citigroup Global Quantitative Research, New York and London.

$$\frac{I}{\sigma} = \gamma\, T\ \mathrm{sgn}(X)\ \left|\frac{X}{VT}\right|^{\alpha} \left(\frac{\Theta}{V}\right)^{\delta} + \langle\text{noise}\rangle$$

$$\frac{1}{\sigma}\left(J - \frac{I}{2}\right) = \eta\ \mathrm{sgn}(X)\ \left|\frac{X}{VT}\right|^{\beta} + \langle\text{noise}\rangle$$

$$\text{Permanent impact:} \quad I = \frac{S_{\text{post}} - S_0}{S_0}$$

$$\text{Realized impact:} \quad J = \frac{\bar{S} - S_0}{S_0}.$$

$$\alpha = 0.891 \pm 0.10$$
$$\delta = 0.267 \pm 0.22$$
$$\beta = 0.600 \pm 0.038.$$

Shares outstanding: We constrain the form of $\mathcal{L}$ to be

$$\mathcal{L} = \left(\frac{\Theta}{V}\right)^{\delta}.$$

where Θ is the total number of shares outstanding, and the exponent δ is to be determined. The dimensionless ratio Θ/V is the inverse of "turnover," the fraction of the company's value traded each day. This is a natural explanatory variable, and has been used in empirical studies such as Breen, Hodrick, and Korajczyk (2002).

Market impact models for trading

Two types of market impact
(both active, both important):

- Permanent
  due to information transmission
  affects public market price
- Temporary
  due to finite instantaneous liquidity
  "private" execution price not reflected in market

Many richer structures are possible

Temporary vs. permanent market impact

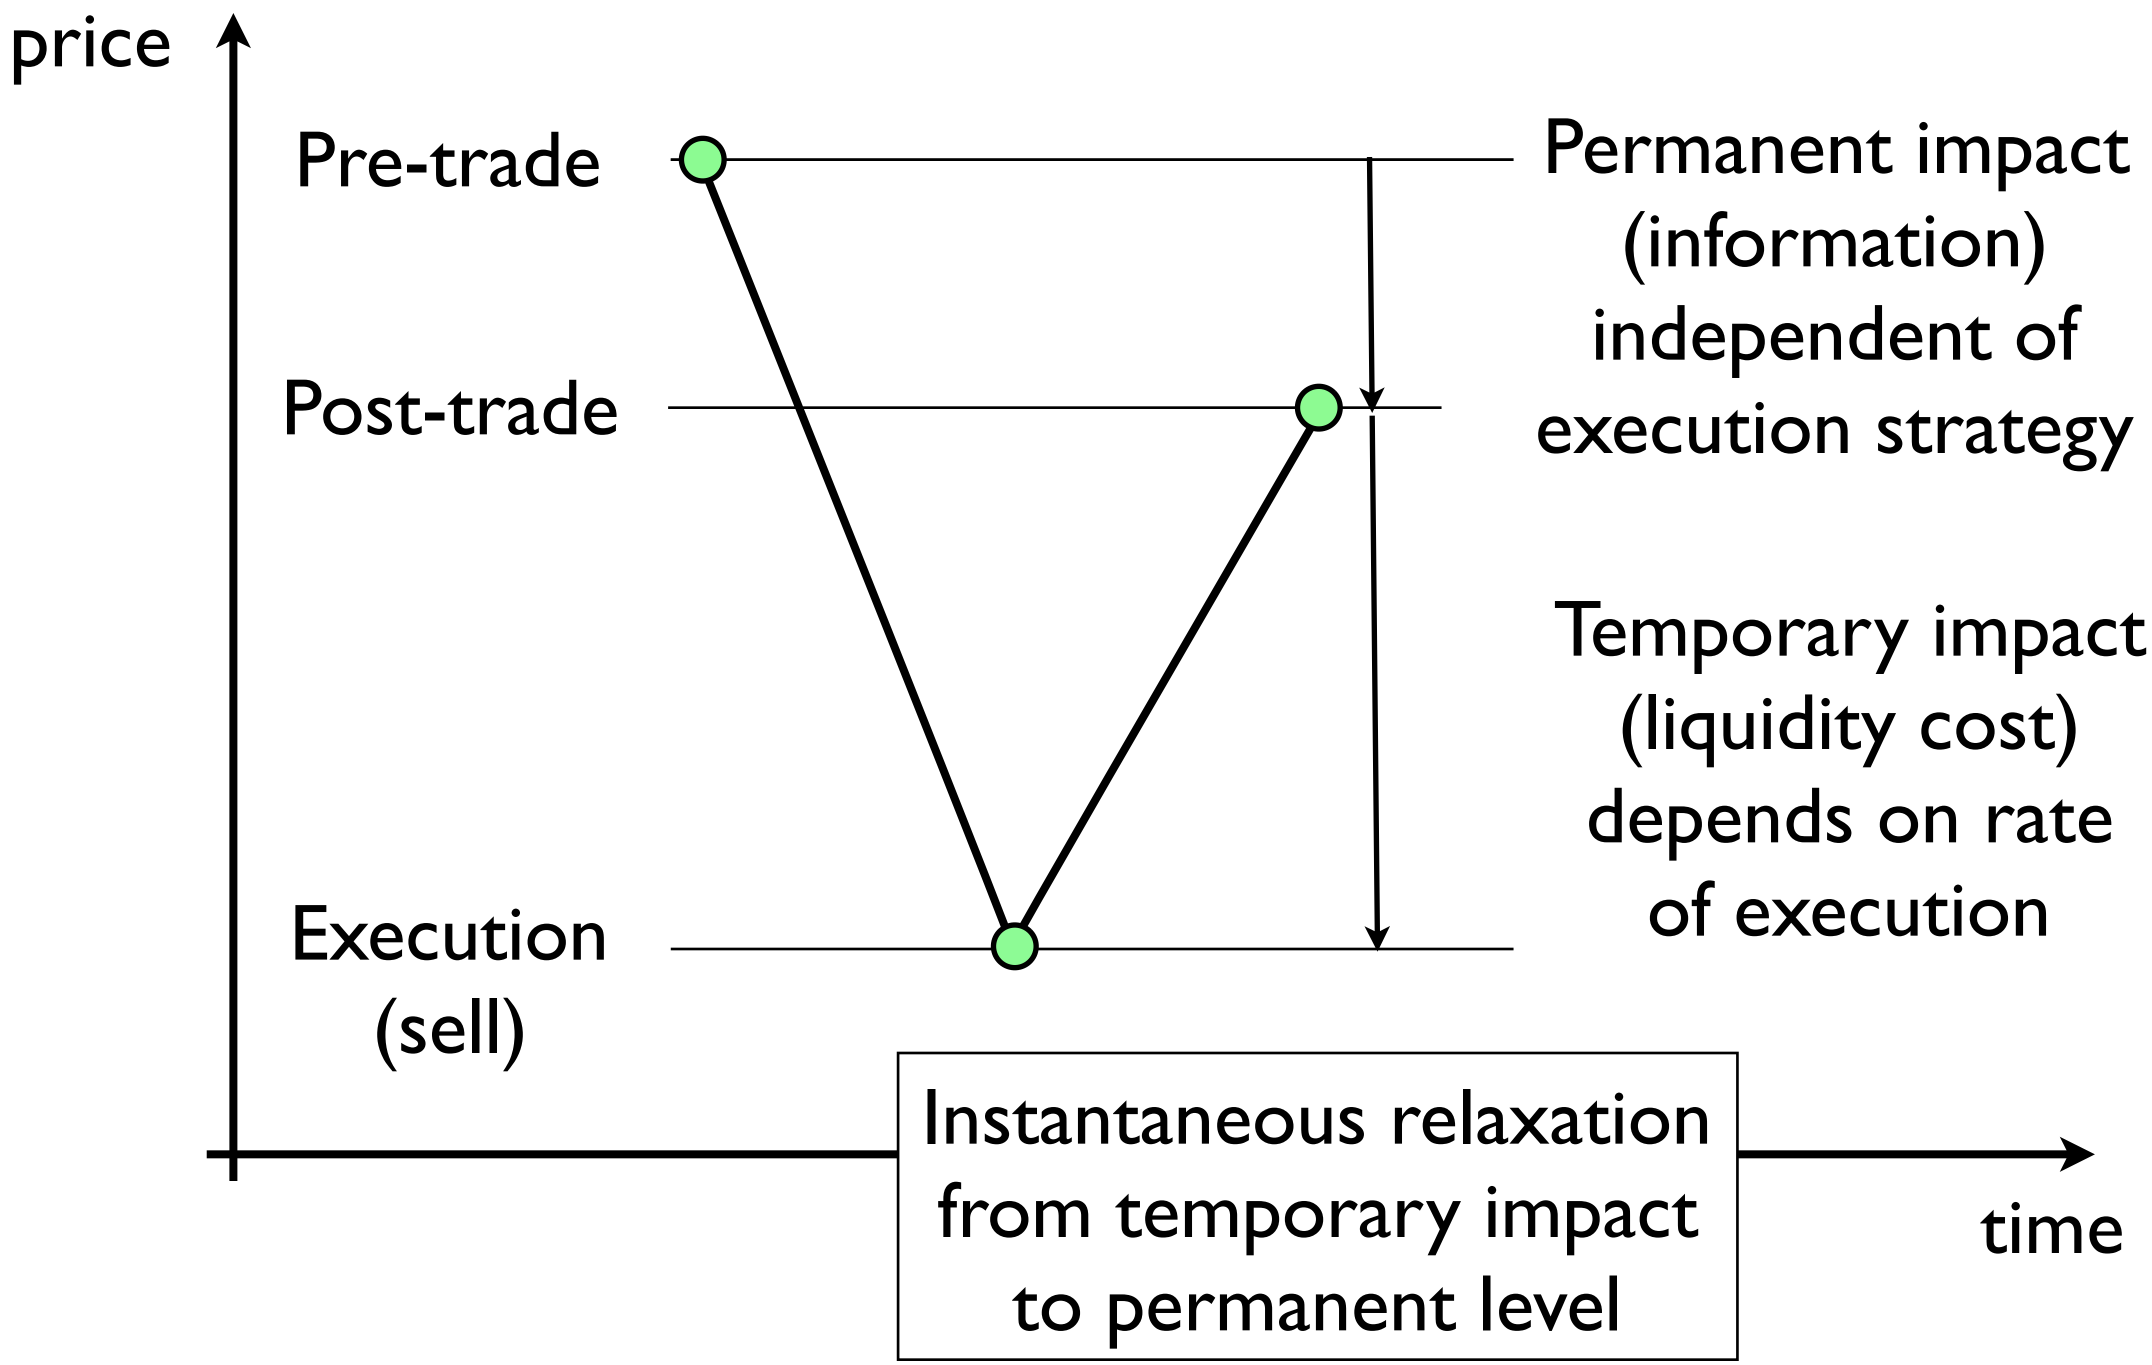

Jim Gatheral: richer time structures for decay

Permanent impact

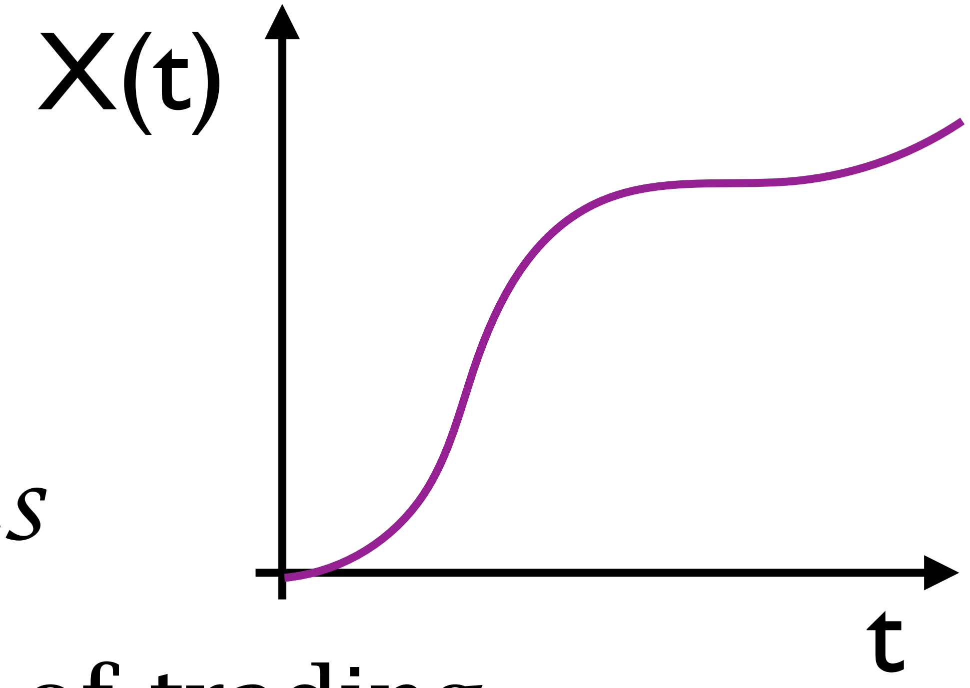

$$X_t = X_0 + \int_0^T \theta_s \, ds$$

$$\theta_t \;=\; \text{instantaneous } \textit{rate} \text{ of trading}$$

$$dP_t \;=\; \sigma \, dW_t \;+\; G(\theta_t) \, dt$$

Linear to avoid round-trip arbitrage (Huberman & Stanzl, Gatheral)

(Schönbucher & Wilmott 2000: knock-out option--also need temporary impact)

$$G(\theta) \;=\; \nu \, \theta$$

$$\boxed{P_t \;=\; P_0 \;+\; \sigma \, W_t \;+\; \nu \, (X_t - X_0)}$$

(independent of path)

Cost to execute net X shares = $\frac{1}{2} \, \nu \, X^2$

Temporary impact

We trade at $\tilde{P}_t \neq P_t$

$\tilde{P}_t$ depends on instantaneous trade rate θ_t

$$\tilde{P}_t = P_t + H(\theta_t)$$

Require finite instantaneous trade rate

$\Rightarrow$ imperfect hedging

Example: bid-ask spread

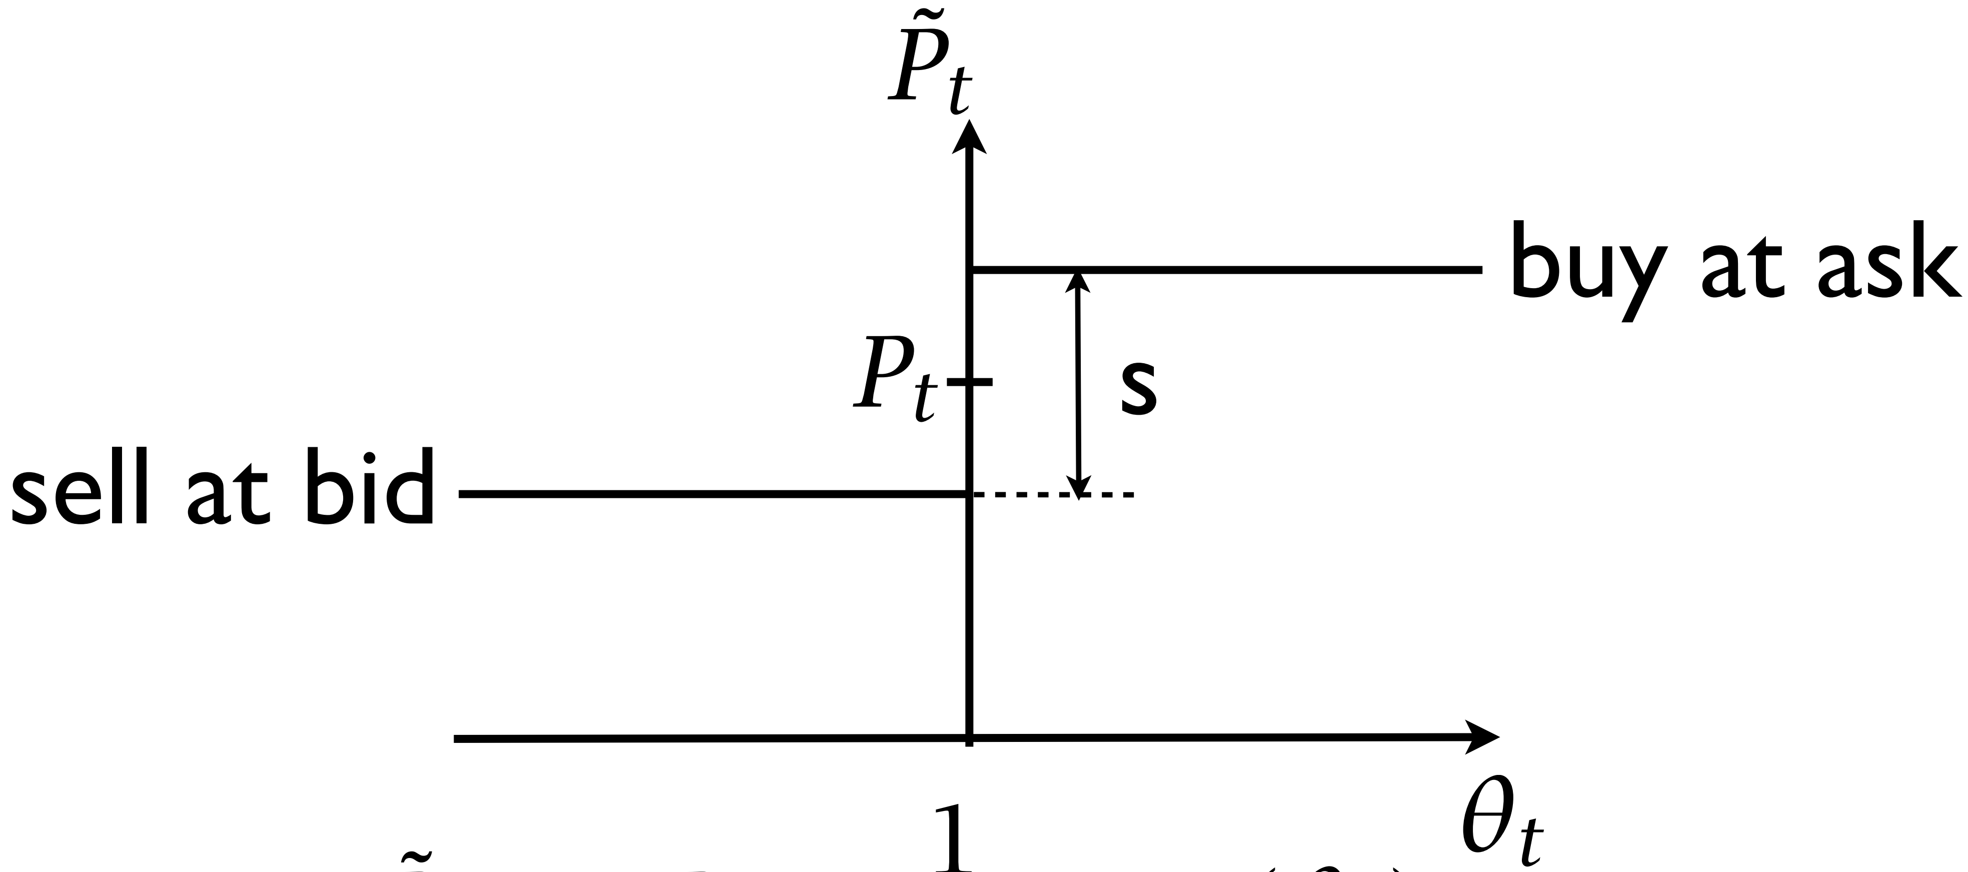

$$\tilde{P}_t \;=\; P_t \;+\; \frac{1}{2}\, s\, \mathrm{sgn}(\theta_t)$$

"Linear" model: cost to trade $\theta_t\,\Delta t$ shares

$$\frac{1}{2}\, s\, \mathrm{sgn}(\theta_t) \cdot \theta_t\,\Delta t \;=\; \frac{1}{2}\, s\, |\theta_t|\, \Delta t$$

Critique of linear cost model

independent of trade size

  not suitable for large traders

in practice, effective execution near midpoint

  spread cost not consistent with modern cost models

  liquidity takers act as liquidity providers

Proportional temporary cost model

$$\tilde{P}_t \;=\; P_t \;+\; H(\theta_t)$$

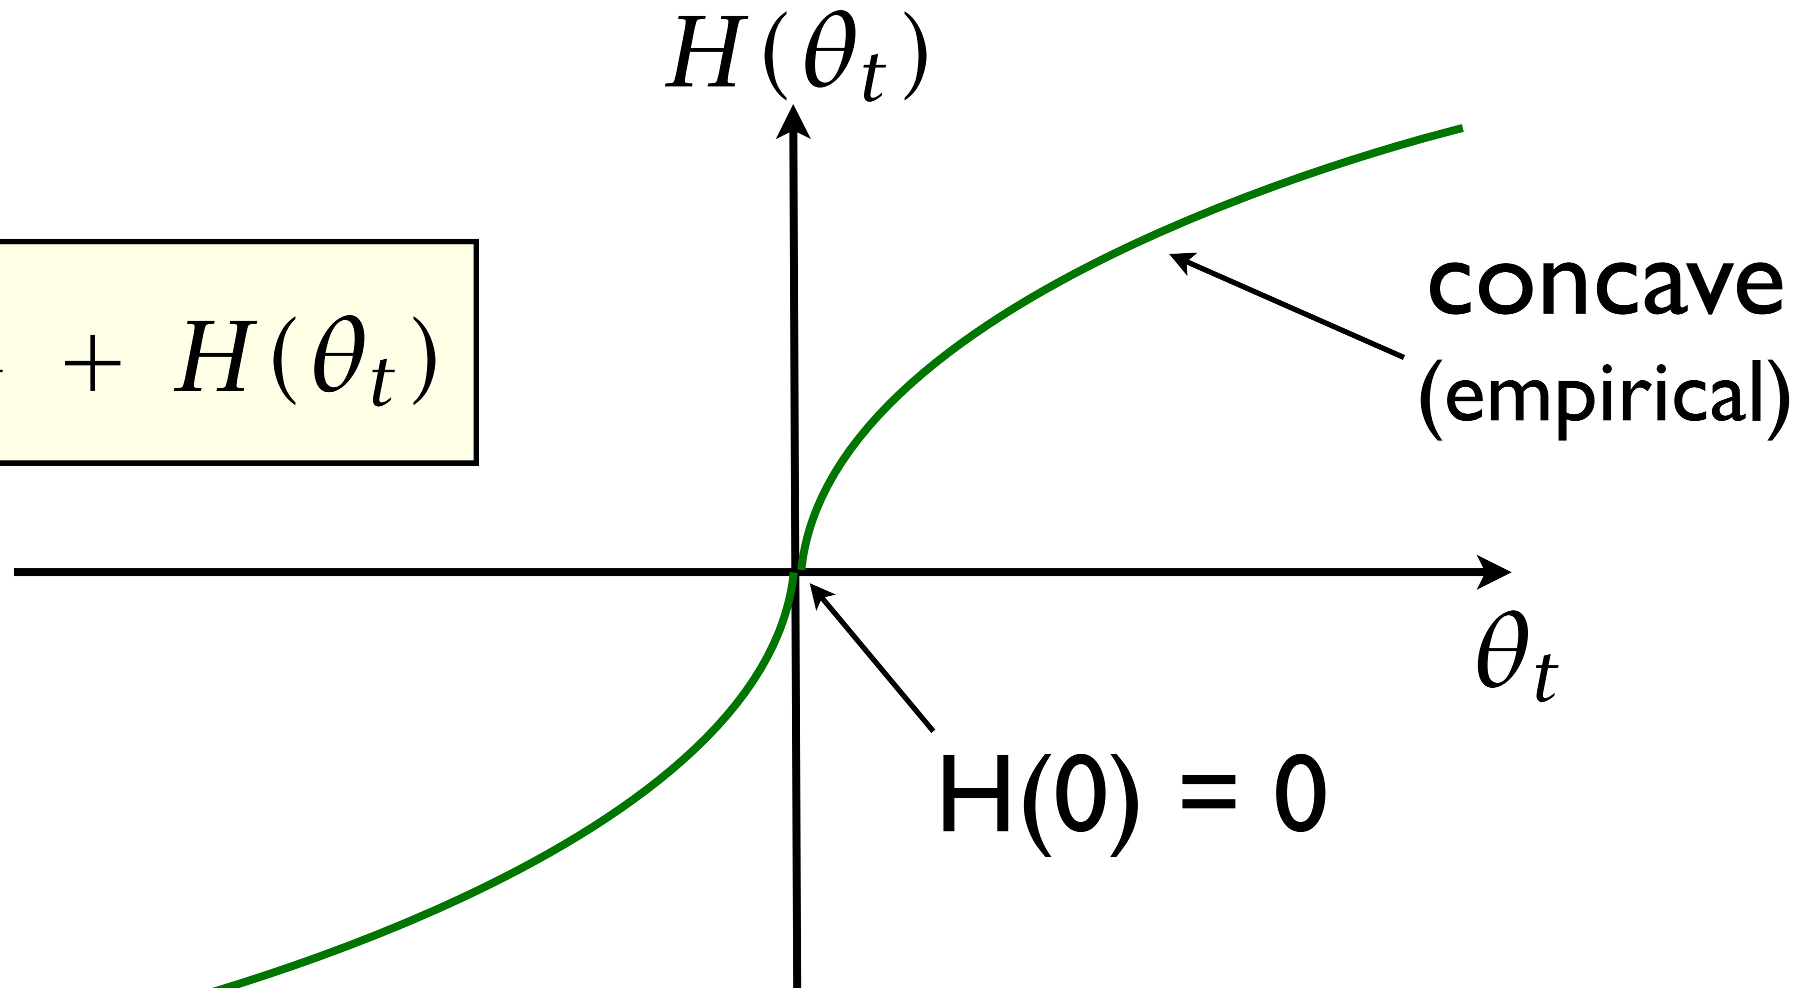

Special case: linear for simplicity $H(\theta) \;=\; \frac{1}{2}\lambda\,\theta$

$\Longrightarrow$ Quadratic total cost: $H(\theta)\cdot\theta\,\Delta t \;=\; \frac{1}{2}\lambda\,\theta^2\,\Delta t$

Conclusions

Market impact not easy to define or measure

- trading and price changes are related
- who pays trading cost to whom

Micro models from public data

- including trade size
- excluding trade size

Macro models from private trade data

- excluding time
- including time